2008

Development of an Inexpensive, Haptic Graphical Display Device

David Burch
Virginia Commonwealth University

School of Engineering

Virginia Commonwealth University

This is to certify that the thesis prepared by David S. Burch entitled *Development of an Inexpensive, Haptic Graphical Display Device* has been approved by his committee as satisfactory completion of the thesis requirement for the degree of Master of Science in Biomedical Engineering.

Dianne Pawluk, Ph.D., Director of Thesis, School of Engineering

---

Ms. Paige Berry, Adjunct Professor, School of Allied Health Professions

---

Ou Bai, Ph.D., School of Engineering

Gerald E. Miller, Ph.D., Chair, Biomedical Engineering, School of Engineering

---

Rosalyn S. Hobson, Ph.D., Associate Dean of Graduate Studies, School of Engineering

---

Russell Jamison, Ph.D., Dean, School of Engineering

F. Douglas Boudinot, Ph.D., Dean, School of Graduate Studies

Date

Development of an Inexpensive, Haptic Graphical Display Device

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

David Spencer Burch
Bachelor of Science, 2006
Department of Biomedical Engineering
Virginia Commonwealth University

Director: Dr. Dianne Pawluk
Associate Professor
Department of Biomedical Engineering

Virginia Commonwealth University
Richmond, Virginia
December, 2008

Acknowledgements

I would like to thank my adviser, Dianne Pawluk, whose thoughtful guidance I would be lost without. Everything I have accomplished in Graduate School has been with her help and I owe her more than I could ever repay.

I would like to thank my parents, Kirby and Donna Burch, for their unrelenting love and support, without whom I would never exist and certainly would not have made it to Graduate School. I would also like to thank my sister, Heather Krauss, and her husband, John Krauss, my brother, Adam Burch, my aunts, Jo Anne Spencer, Judith Miller, and Jane Gross for all of their love and support as well.

I would like to thank my friends, especially my fellow biomedical engineers Pallavi Ramnarain and Isti Arief, and my haptics labmates, Ravi Rastogi, Justin Owen, Julie Petro, Devnath Vasudevan, and Ketan Vidwans. Victory shall be ours!

I would also like to thank Paul Wetzel, Pete Robinson, Susan Lederman, Paige Berry, Luke Hargrove and Gary Bowlin for their professional opinions, and Joyce Wilkins for her more than professional opinion, about anything I wanted to know.

Table of Contents

Abstract

DEVELOPMENT OF AN INEXPENSIVE, HAPTIC GRAPHICAL DISPLAY DEVICE

By: David Burch

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

Virginia Commonwealth University, 2008

Director: Dr. Dianne Pawluk, Associate Professor, Biomedical Engineering

A finger-worn haptic device capable of rendering 2-D graphics through vibrotactile feedback is presented. The device development is presented from its initial stages of being a stylus design using a photo-interrupter optical sensor and pager-motor actuator to a small case worn on the finger using a RGB color sensor and a piezoelectric actuator. Testing of the latest prototype design shows that it has a spatial sensitivity (<2mm) comparable to natural touch (~1mm) and can be used to output a variety of vibrotactile textures. The design can be expanded for a multiple finger, independent device, while remaining affordable (<$100) and highly portable (<500g).

1. Introduction

For individuals with sight, graphical visual representations are a universal means for conveying *unfamiliar* information. Their use ranges from teaching young children some of their first words, to guiding tourists without concern for language barriers, to helping visualize complex data. However, for individuals who are visually impaired, there is a dearth of appropriate tools and representations to obtain access to this same information. An affordable and commonly used alternative, to *replace* a graphic using words in text or auditory form, can be very useful, cost-effective and significantly more accessible. However, there are many situations for which words, either in text form or speech, are simply inadequate. Words cannot be used in situations where language barriers exist, regardless of form. Nor are they appropriate when teaching children language, as the visuals act as a crucial supplement to instruction. While word descriptions are an important facet of learning, they alone are no replacement for being able to *independently discover* patterns and spatial relationships in information. For example, displaying time-series data and its analyses in a graph to look for spatial patterns is a fundamental way of enhancing insight into a scientific experiment or financial situation. Determining spatial relationships can also be particularly important for understanding how machinery and devices should be used in the workplace, as well as the spatial layout of a person's work environment. Having access to all these types of information easily would allow a person who is visually impaired to perform more tasks independently, improving both their self-esteem and value in the workplace.

To mitigate the limitations of descriptive language, visual information can be rendered in a form that utilizes an alternative sensory system; the most common alternatives are the auditory or haptic system. The auditory system is a well-known and studied system; however, its use precludes accessibility by those who are deaf-blind. Also, Jeong and Gluck (2002) noted that during map feature identification tests, subjects using auditory feedback (sound volume), haptic feedback (extent of vibration), or both performed best when using only haptic feedback [40]. SoundView, a system designed by Kees van den Doel [41], applies sonification in a similar manner to the device described in this thesis: image colors are translated to associated "roughness" encoded by varying scrapping sounds. Specifically, they encode such color characteristics as hue, saturation, and brightness by altering the digital filter characteristics for the scrapping sound output; their results. However, the use of non-speech auditory feedback as a substitute for visual feedback can interfere with speech recognition due to masking effects. Such auditory masking can inhibit learning during classroom instruction where normally visual and auditory information are present simultaneously. In addition, hearing has no correlate to using multiple fingers, a potential method to speed up the very slow, serial processing of information that occurs by audition and haptics with one finger. Thus, the haptic system, a combination of the tactile and kinesthetic senses, is an alternative system more widely accessible.

For haptics, directly copying a visual image using means of relief intensity mapping of the image is known to be extremely ineffective [1]. Therefore, the most frequent haptic method of representation is the use of static raised-line drawings: simplified

outline drawings where the lines are permanently elevated above the background on a piece of paper. After attending several focus groups with educators who specialize in teaching those with visual impairments, a common theme crystallized: that the current methods of conveying visual information are too inefficient in terms of cost or time, and too ineffective at helping people in the task. These educators and assistants proved much anecdotal evidence of having to spend all night to create tactile versions for all the visual graphics for a science or math course for just one day. The methods used to mass produce these tactile renderings vary greatly, creating different perceptions of a single image produced by different methods. Depending on the resources available, a school system will only have enough funds to invest in a single method of tactile image creation system. Unfortunately, the method chosen is not necessarily the method used on standardized tests, creating confusion and poorer results among students who are blind or visually impaired. Some educators for these students have given up on using tactile graphics, claiming they are too ineffective and too expensive to justify. Others have found success in creating tactile experience pictures made of different materials, but will readily admit the method takes far too long.

Momentarily ignoring the issues of cost and ineffectiveness of static raised line images, the time required to produce a single unique tactile graphic is burdensome enough. From the experience of producing static tactile images in the lab, a single image can take anywhere between twenty minutes to over an hour to produce. Reproduction of the same image (making copies) can take approximately two to five minutes per copy, and, depending on the method, has a high rate of error (paper jams, smearing or bleeding,

et cetera). These problems only become compounded when an individual needs to access a large set of images, such as images of biological specimens or graphs of data. To produce such a set of static tactile images, someone would have to spend hours producing them; however, a method that could dynamically represent the images could reduce the down time significantly. Dynamic representation involves some sort of device that can temporally present different media (in this case, a tactile image), without needing to replace components, and typically involves some level of user control over what is currently being presented.

The goal of this project was to develop a means to render 2-D graphics haptically, in a cost-effective, highly usable means that has a higher degree of accuracy than current raised-line images. Such a system had to be easily usable in a highly interactive environment, such as when performing data analysis or navigating the Web. This means that the display, itself, needed to be easily and frequently refreshable. Second, as many individuals who are blind or visually impaired live near or below the poverty line, cost was a critical issue: not only in terms of "maintenance" cost but also the initial cost of the device. A low cost device was considered crucial for acceptance by the population of people who are blind or visually impaired. Other criteria included that the device be: (a) easily portable, (b) easy to use and maintain, (c) comfortable and safe to use and (d) usable by all segments of the population of people who are blind or visually impaired, with particular consideration for individuals who are deaf-blind and lack some or all form of hearing. By portable, the device should be small in mass, so that a child can easily carry it: 2kg or less. People should be able to use the device in a diversity range of

settings: in the office, at the school, at home, even in outside environments. This dictates that the device should have an internal power supply, like a battery. The device must be easy to use and maintain: it should not require a complicated or lengthy user's manual in order to operate. People should be able to intuitively learn how to operate the device. It should be easy to repair by moderately skilled technician if parts fail without having to be shipped off to repair. A dynamic tactile device that meets these criteria, with the additional ability to produce textures, is presented here; it should be noted that the use of even a single solid "texture" surface has been shown to be significantly more effective in conveying information about an object than raised line drawings [2].

In this thesis the haptic system and its characteristics will first be discussed, with emphasis placed on the elements that influence the design of the device, including spatial resolution, vibration sensitivity, and texture sensing. Next, current methods of rendering tactile graphics using both static and dynamic means will be discussed; devices capable of dynamic rendering will be broken down into point-contact displays and distributed displays. Then, I will discuss the prototype device, starting with the earliest devices (the stylus and the glove), and ending with the current design (the finger-worn model). All the prototype testing performed will be discussed in the same section as the associated prototype model. The devices will be assessed as to how well they fit the overall device design criteria of usability, accessibility, and user safety.

2. Background

2.1: Haptic System Overview

The Haptic system consists of two sensory systems: the tactile sense and the kinesthetic sense. The tactile sense is composed of various cutaneous mechanoreceptors sensitive to various forms of physical contact, nociceptors sensitive to pain, and temperature sensors. Of primary concern are the four types of cutaneous mechanoreceptors, which can be discerned by their terminal structures; they are the Pacinian corpuscle, the Meissner's corpuscle, Merkel's disks, and Ruffini's endings. All of these receptors are long receptors, which can terminate in multiple end organs, but the signal originates from the end organs and does not have an additional connection to the receptor fiber. The afferent fibers for these receptors can be described by their associative receptor field size (small—Type I or large—Type II) and whether they slowly or rapidly adapt to stimuli [3]. The Slowly Adapting Type I receptors (SA I) are sensitive to static pressures (although they respond at higher frequencies as well) and have small receptive field sizes (~2-3mmφ). The Slowly Adapting Type II receptors (SA II) are associated with sensing skin stretch. The Fast Adapting Type I and Type II receptors both detect border transitions and vibrations. All the afferent fibers are myelinated fibers, with all the fiber diameters large, except for those of the Merkel disks, which are small. This combined with long receptor-type contributes to fast conduction times to the brain, on the order of tens of milliseconds.

The kinesthetic sense is compromised of three types of mechanoreceptors: the primary and secondary muscle spindles and the Golgi Tendon Organs. These afferent neurons

send information about muscle stretch and force to the central nervous system, which determines current body position and tracks changes to predict and accommodate for changes in muscle force necessary to maintain position and prevent hyperextension. The muscle spindles are found in groups of specialized fine muscles fibers, called intrafusal muscle fibers; each spindle contains two bag fibers and about five chain fibers. The nuclear bag fibers can be classified as either static or dynamic fibers and are surrounded by annulospiral rings of two types of afferents, the Ia primary afferents and the II afferents. Both the nuclear bag fibers and the chain fibers contain Ia primary afferents, but only the static nuclear bag fibers and the nuclear chain fibers contain II afferents. The Ia primary afferents are rapidly adapting, large myelinated nerve fibers that surround the central, non-contractile portion of the muscle fibers and have stretch-activated sodium channels that correspondingly increase the neuron firing rate to the rate of muscle stretch. The II afferents are slowly adapting, medium myelinated fibers found between regions of the fibers and respond to static muscle length. The sensitivity of these receptors is maintained by a gamma motor efferent, which controls the length of the intrafusal muscle fibers. Muscle spindle density depends on what type of movement that muscle performs; muscles which are part of fine movement motor groups, such as those in the fingers, may contain hundreds of spindles, while larger motor groups responsible for gross movement, such as those in the upper arm, may have only one per muscle fiber. However, the higher concentration of spindles in fine movement motor groups does not necessarily correspond to higher spatial acuity as they do with vision or touch [3]. Golgi tendon organs are receptors found within the collagen fibers of tendons, in series to the skeletal muscle, and

in ligaments. Like muscle spindles afferents, Golgi tendon organs have stretch-activated sodium channels; however, because they are in series with the muscle the stretch they sense corresponds to the force developed by the muscle. The Ib afferents carry signals from the Golgi tendon organs to the spinal cord, where through interneurons they make connections to alpha motor neurons, inhibiting their firing and relaxing the muscle.

2.1.1: Spatial resolution

The spatial resolution of the haptic system depends on various factors and thus, can be expressed in various ways. First, the mechanoreceptors that make up the haptic sensory system have independent receptive field sizes that limit their individual resolution. The SA I receptors are highly dense (1 per $1mm^2$) complexes, which are sensitive to fine spatial details rather than overall geometric shape. Their disk-like structure makes them robust and fairly invariant in their response to contact forces ranging from 0.2 to 1N [4], while their spatial resolution is fairly consistent for scanning velocities up to 80mm/s [5]. The FA I (or RA) receptors are the densest receptor types in the fingertips, with an average density of 1.5 receptors per mm^2 [3]. Like the SA I receptors, they have a narrowly defined receptive field around the same size of a few millimeters in diameter, but because they respond evenly across their receptive field, they have poorer fine detail resolution than the SA I receptors. Both the SA II afferents and the FA II afferents are far less populous than either the SA I or FA I type receptors, and have notably larger receptive field sizes, extending over several square centimeters.

When evaluating the prototype devices, common testing methods of determining spatial resolution will be employed. The two-point discrimination method is one of the

most common methods, with normal threshold values being 2-4mm on the fingertips and 10-11 mm on the palm [5]. However, the two-point discrimination test has variability, as different contact forces can be applied. Another method is a two-case forced choice method where subjects must determine the orientation (horizontal or vertical) of rows of raised lines. The test involves presenting fixed patterns of either horizontally or vertically orientated raised lines or groves for a number of spatial periods. Subjects must determine whether the presentation is of vertical or horizontal lines for each stimulus. Using this method a more accurate spatial acuity of around 1mm was observed [6]. Yet another method is point localization, which is presented as a two-point discrimination test with the presentation of the two points temporally separated. The subject must decide whether the two points were at the same or different locations. Resolutions found with this method tend to be in-between two-point discrimination and two-alternative forced-choice testing, typically around 1 to 2mm [5]. Hyperacuity, or spatial discrimination thresholds finer than typical resolution tests yield, has been found to exist in tactile sensing [43]. Loomis (1980) found that in some tasks, people exhibit thresholds much lower than the standard two-point discrimination task, with some tasks yielding discrimination thresholds at around 200μm. These resolutions are smaller than the receptive field sizes for the mechanoreceptors would naturally allow for, and might be due to the relative excitation of neighboring mechanoreceptors causing lateral inhibition. This would suggest that the hyperacuity thresholds are limited only by the spatial precision of the elements causing the inhibition. However, this does not seem to be the case for touch, as the point localization thresholds were lower than for the other hyperacuity configurations.

Thus, it is believed that hyperacuity results not from a response to spatial frequency, but rather spatial phase.

The discrimination between tactile lines of different thicknesses shows that as the thickness of the line increases, the threshold for detecting differences in the lines decreases, to a point [24]. Berla and Murr (1975) tested subjects on raised line stimuli that had five standard line thicknesses and thirteen variable line widths (e.g. the 1.02mm standard had 13 variations to it, in 0.102mm increments from the standard size). Subjects were presented one of the standards along with one of the variables and were asked to judge which one had thicker lines. They noticed that the Weber fractions decreased exponentially for raised-lines up to 6.35 mm, but then increased for raised-lines with a 12.7 mm thickness, indicating a U-shaped sensitivity response [24].

2.1.2: Vibration Sensitivity

The human sense of touch can perceive a range of vibrotactile stimuli with the ability to discern changes in both amplitude and frequency of the stimuli [51].The threshold sensitivity to vibration stimuli generally follows a U-shaped curve (as it does for many psychophysical thresholds, including lines of constant magnitude), with the lowest thresholds at frequencies around 250 Hz [3]. The theory behind vibrotactile sensing is that there are four channels of perceptual processing, each sensitive to different constitutive characteristics. The four channels can be broken down into the PC channel and three non-PC (NPI-NPIII) channels. The PC channel is, as one would expect, consists of the FAII mechanoreceptors; they are sensitive to frequencies between 35-500 Hz with peak sensitivity around 250 Hz, ±50 Hz. The NPI channel is responsive to frequencies

ranging from 3-100 Hz for small contactors (0.008 cm^2) and 3-35 Hz for larger contactors (2.9 cm^2). Unlike the PC channel, the NPI channel has a rather flat frequency response, with a slight improvement around 30 Hz, ±5 Hz; it is supposed that the NPI channel consists of the FAI afferents. The NPII channel is responsive to frequencies ranging from 80-500 Hz for small contactors, with similar frequency response and peak sensitivity as the PC channel; it is suspected that the NPII channel consists of the SAII afferents, though most of their characteristics other than contactor size are similar to the PC channel. The NPIII channel is responsive to very low frequencies, typically around 0.4-3 Hz regardless of contactor size, and has a flat frequency response with peak sensitive at the lower end of their range of sensitivity; this channel is suspected to consist of SAI afferents [3].

Psychophysical results related to vibration as an encoding method found that subjects tested with variable frequencies and amplitudes had longer reaction times for increases to both frequency and amplitude, with a greater correlation between the increases in reaction time and the increases in amplitude [51]. Additionally, the experiment exhibited significant interaction between the two dimensions of amplitude and frequency, such that variation in one dimension was correlated to proportional changes in the perception of the other dimension. This result reduces the effectiveness of using both amplitude and frequency to encode information with vibrotactile feedback—an important consideration for the design of the device.

It is interesting to note that non-PC channels (lower frequency vibrations) show no temporal or spatial summation; an important factor because it means that these channels have the same stimulation thresholds regardless of the area of contact or the duration of the vibration stimulus. This supports the paradigm that different channels of vibration sensing are associated with different mechanoreceptors. Further, selective fatigue only occurs within the same channel and not across multiple channels; also selective enhancement due to temporal or spatial summation only occurs within the PC channel [44]. Temporal masking effects are also selective (i.e. only observed within channels, but not across them), illustrating again that multiple mechanoreceptors are responsible for sensing vibration. Another issue is vibration masking, which is similar to adaptation in that it affects responses within a single channel. Temporal masking can occur when a stimulus with a lower threshold is temporally collocated near another signal (e.g. a 250Hz vibration is presented temporally near a 400Hz vibration); the closer the masking stimulus is presented to the signal stimulus, the greater the masking of the less-sensitive vibration. Temporal integration of complex vibrotactile signals compromised of multi-tone (or frequency) notes can increase errors in perception, as some components might mask others. Simpler notes containing less components show a greater accuracy in perception [25]. Tan and her colleagues have found [26] that the Weber fractions for signals presented along with maskers will decrease as the signal amplitude increases, and amplitude discrimination thresholds are higher for lower amplitude signals presented with maskers.

Additionally, all channels for vibration perception are prone to adaptation with long duration of exposures. For all channels, the frequencies that the channel is most sensitive to are also most likely to cause adaptation, or response fatigue. Fatigue causes an increase in the thresholds, desensitizing the mechanoreceptors to stimulation. This is a deleterious effect for a device for individuals who are blind as it compounds the difficulty of using the haptic system. The factors of vibration threshold sensing, adaptation, temporal summation, and masking play an important role in choosing what vibrotactile stimuli to use with the device. Vibration frequencies chosen should not be in the range of those with the greatest rate of adaptation (i.e. the most sensitive frequencies) for any channels to reduce stimulus fatigue. Stimuli involving multiple frequencies ideally should be comprised of frequencies sensed by different perceptual channels to avoid possible masking affects.

2.1.3: Texture Sensing

The perception of surface texture is both a multimodal and multidimensional sensation; it is a combination of both mechanoreceptor stimulation and surface characteristics. The surface characteristics most often used to describe texture are the apparent roughness/smoothness, hardness/softness, and a third descriptive dimension that is harder to classify [3]. Hollins and colleagues (2000) described this third dimension as sticky/slippery; however, a more appropriate surface characteristic might be adhesiveness/cohesiveness (a glue versus a lubricant) [21]. An important characteristic of texture sensing is how accurate people are at doing it; people are capable of

discriminating sub-micron deviations [22] and minor deviations (2%) in spatial periods in texture patterns with 75% accuracy [23].

The primary focus on texture perception has been the characteristic of surface roughness. Research has shown [10, 23] that the perception of roughness is independent of active haptic sensing and determined through cutaneous sensing. The characteristics of roughness can be defined by two variables: interaction with the surface and surface constitution. Surface roughness is dependent on the user's interaction with the surface; primarily, the force of contact and the relative speed between the surface and skin help determine the perception of roughness. The surface constitution also plays a significant role: specifically the spatial period of the elements and their spatial duty cycle (percent element width or percent inter-element width versus the spatial period). The current theory of roughness sensing suggests that roughness is encoded separately for coarse roughness (gratings with spatial periods >200μm) and fine roughness (gratings with spatial periods <200μm). Coarse roughness is intensively coded through skin deformation, but often includes temporally-encoded (i.e. vibration) elements relating to sharp transitions. However, the temporal components of coarse roughness do not greatly influence the overall perception of roughness, and therefore coarse roughness is not prone to adaptation. Vibration sensing plays an insignificant role in the perception of coarse roughness [10, 45], except in some cases of indirect contact with the surface, such as when a subject uses a rigid probe to sense the surface [45]. Furthermore, the spatial period does not influence the perception of roughness as much as the inter-element spacing, which shows a linear trend between spacing and perceived roughness [3]. For

fine roughness, vibration sensing appears to be the only mode of perception, and textures with spatial periods smaller than 200μm are affected by adapting stimuli [45, 46], and are similar for both direct touch and indirect touch [45].

2.1.4: Cognitive Processing of Haptic Feedback

The main concepts of haptic processing important to the design considerations for the device are: surface properties identification, object recognition, and vibrotactile signal apperception. When people touch an object, they can immediately encode surface properties about the object; within a few seconds of active exploration, people build a mental representation of the object, enabling them to identify the object with remarkable accuracy. This effect appears to be the result of parallel processing within the haptic system, which is the ability to simultaneously process haptic feedback across multiple fingers. Lederman and Klatzky have shown that many coarse-intensive discriminations are processed early on, but that spatial dimensions (i.e. "a domain of variation that is accessible to the perceptual system") are processed later on [36]. Of the early-stage discriminations, material properties such as surface texture, thermal conductance, and relative hardness all exhibit signs of parallel processing (i.e. as described in [36] having flat or shallow search functions). Abrupt surface discontinuities and rough 3-D geometric shape also exhibit signs of parallel processing, but contour following does not. This means material properties, such as surface roughness, can be intensively coded across multiple fingers in a relatively short interval of around 200ms. This is in marked contrast to the encoding of fine spatial discriminations, as is typically performed in a contour following task, which are serially processed over much longer intervals.

However, it is important to note that only coarse (as opposed to detailed) haptic information is processed in parallel and that even parallel processing does not necessarily ensure accuracy. Craig [47] addressed whether vibrotactile spatial patterns generated by a vibrating pin array distributed over two fingers (index/middle fingers or both index fingers) affect haptic perception. The 144 pin array was distributed over 11mm x 27mm on the fingerpad and set to vibrate at 230Hz; this is similar to the Optacon device which is discussed later in the thesis. Subjects were asked to perform a variety of discrimination and identification tasks; patterns would be presented asynchronously to the two fingers, with the stimulus onset asynchrony being varied. Craig found that attention was best when the stimulus was presented to only one finger and notably worse when present to both the index and middle fingers on the same hand. Stimulus discrimination and identification also improved with increasing asynchrony. In contrast, when the stimuli were presented bilaterally (across two hands), discrimination percent correct was higher than ipsilateral percent correct, and was relatively unaffected by increasing the asynchrony. However, Craig suggests that masking effects could have influenced the ipsilateral task performance. Later [48], Craig and Qian found that temporal interference and masking effects functions were similar for one and two finger patterns. Additionally, their results suggest that dividing attention between two bilateral fingers does not alter any temporal interference, and therefore, should not inhibit pattern processing in bilateral exploration.

The final consideration regarding haptic cognition for the developed device was the apperception of vibrotactile signals. Vibrotactile signals when given associated meaning

can convey non-visual information. Often these signals are referred to as tactons, or tactile icons [49]; however, the term tacton has typically been used to mean a haptic icon, or haptic reference point. For the purposes of this research project, I use the term texture rather than tacton, since the device will encode information about an area, not a point. Textures can encode information by varying the parameters of vibrotactile stimuli, such as frequency, amplitude, duration, waveform, summation, and modulation; however, some of these parameters are of limited use [50] and others are device dependent (i.e. the ability of the device to output a range of frequencies or amplitudes may be mechanically limited). Further, given the interaction effects with the perception of both amplitude and frequency [51], it is best to use one or the other, with frequency having the wider range of discrimination levels than amplitude. Waveform [52] and frequency [52, 53] have both been shown to be fairly good parameters to use to create sets of icons; however, some preliminary testing performed in the lab suggests that there are limits to their combined usefulness.

Studies have shown that modulated vibratory signals (250 Hz carrier with 20, 30, 40, and 50 Hz modulators) and unmodulated (250Hz sine wave) signals have significantly different perceptions, but that between levels of modulation there is reduced discrimination [49]. Weisenberger [42] found that higher frequency carriers (50, 250 Hz) were identified correctly with greater frequency than a lower frequency carrier (25Hz) over a range of modulation frequencies, with the 15 to 30 Hz range having the best performance with the two higher frequency carriers. Performance was notably poor when the carrier and modulator frequencies were similar; additionally, the 100 Hz carrier had

overall much poorer performance than any of the other carriers. Weisenberger suggests that this result may be because 100 Hz lies within the range of both the PC and non-PC channels, and, therefore, the sidebands of the modulated signal might stimulate both mechanoreceptive systems, causing signal competition.

2.2: Tactile Imaging Methods

2.2.1: Static Images

Static tactile images are the dominant form for rendering 2-D graphical information to individuals who are blind or have low vision. The most common method is to create a relief copy of the image (or in many cases, a very simplified version of the image). The methods to create the relief images range from raised-line drawings produced with thermal-form swell paper to vacuum-form "structures" [11]. Raised-line drawings require the use of a unique paper embedded with alcohol-filled microcapsules; these microcapsules become 'activated' when coated with ink or certain graphite waxes, causing the microcapsules within the colored portions of the paper to burst when heated in a thermal-form 'tactile enhancer'. (Some paper types require an ink to work, but others seem to work when coated with a 'dark' substance that absorbs more heat. These effects typically require different paper or higher settings on the thermal-forming heater.) Raised-line images only have one "level" of output; either the microcapsules burst causing the paper to swell-up, or they don't. However, the bursting of the microcapsules is often very inconsistent, causing deviations in the height of the swelling. To minimize any inaccurate perception of multiple levels caused by these deviations, raised-line drawings are restricted to only using lines (hence the name) to depict graphics: solid

objects must be converted to outline representations. Magee and Kennedy [54] found that active exploration of raised-line drawings to have very poor performance; this conclusion was reaffirmed by the anecdotal evidence provided by educational coordinators for children who are blind or have visual impairments. Passive, or guided exploration, does increase identification [54, 55]; however, it also reduces user independence since it requires a sighted individual to guide the user's hand.

Other means to produce tactile images are through using embossing printers or vacuum-formers, which physically deform a material, and tactile experience pictures, which are images constructed from a variety of materials with different surface properties [11]. Both embossers and vacuum-formers tend to require very expensive equipment in order to make, but can produce tactile images with multiple topological levels. (That is, control the depth of deformation, either at discrete levels, such as the Tiger Embosser, or continuous depths, such as vacuum-forming machines.) Tactile experience images tend to have anecdotally very high performance rates, perhaps due to individuals' ability to haptically discriminate different textures and process them in parallel across multiple fingers. However, these images take a long time to make, with each successive copy taking as long as the initial image to make (non-scalability of production).

Nevertheless, Thompson and her colleagues have developed a means to redesign tactile pictures [2, 7] based on some of the work of Kurze [56]. Their method, called TexyForm, involved encoding 3-D object orientation into a 2-D representation of the object. For individuals trying to use touch to obtain a linear representation of a 3-D

object, parts should not be occluded as sighted individuals might draw them. Also, the shape and size of an object or part does not change with the direction from which it is explored [7]. Figure 1a shows an example of TexyForm patterns, as printed in [7]. TexyForm uses coarse textures, such as the three images on the left of Figure 1a, to represent 3-D orientation, and uses fine textures, such as the three images on the right of Figure 1a, to separate object parts. Figure 1b shows actually TexyForm test images and how the two texture types are used to render an image [2]. They found that using TexyForm compared to visually realistic images, that early-blind subjects had 50% accuracy (versus 12.5% for visually realistic) and late-blind subjects had near 70% accuracy (versus 44%), with some subjects having 100% accuracy. The smallest gain was with blindfolded sighted subjects, with a 56% accuracy (versus a 50% accuracy for visually realistic) [7].

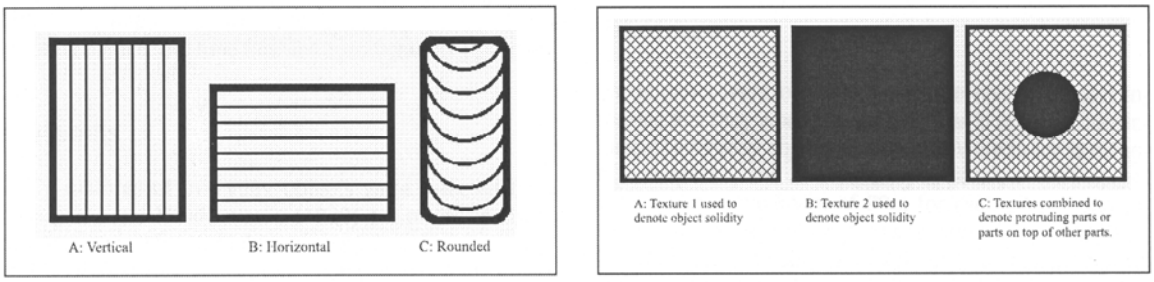

Figure 1a TexyForm patterns [7]

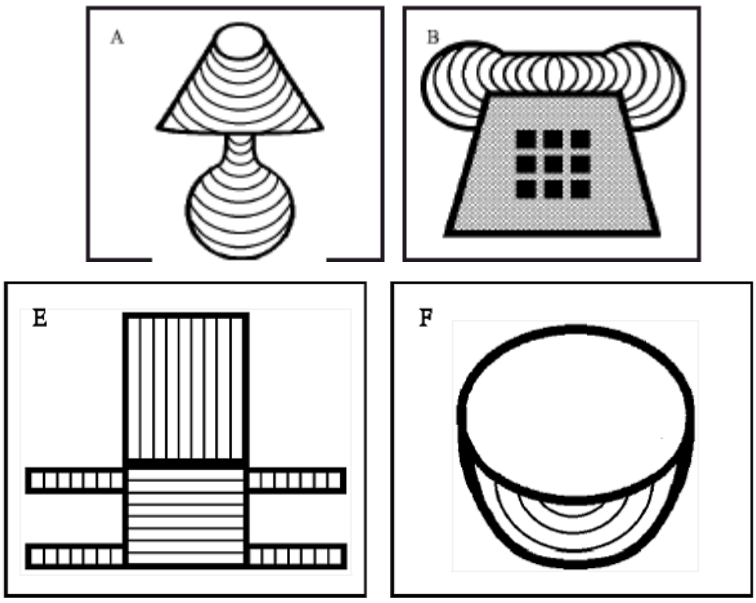

Figure 1b Actual TexyForm Test Images [2]

2.2.2: Dynamic Display Methods

In addition to the static methods of representing tactile graphics mentioned previously, there have been several dynamic displays developed or proposed (for recent reviews, see [3, 17, 18]). These displays can be divided into two main categories: point contact displays and distributed tactile displays.

2.2.2.1: Point-Contact Displays

Point contact displays are displays that model the contact between the graphic and the user (whether the finger or whole hand) as a point interaction (in contrast to direct contact with our fingers where we feel a contiguous, spatially-distributed representation across each fingertip). Appropriate feedback to describe the graphic is given based on the location of the point of interaction. These systems include those that represent boundaries of objects or graphics by nonspeech sounds (e.g., Kurze [19]), forces (e.g., the

PHANToM force-feedback device, SensAble Technologies, Inc. Cambridge, MA and the Wingman Force-feedback Mouse, Logitech) or vibration (e.g., Kurze [20]).

However, there are several problems with all of these devices. On a basic level, the use of sonification (nonspeech sounds) is problematic in that, although cost effective, it precludes access for deaf/blind individuals. The most commonly proposed force-feedback device to be used is the PHANToM [27-29], which is very expensive and the benefits not obvious [28, 29]. The use of more affordable force-feedback mice [30-32] has been proposed; however, their workspace and the forces generated are very small. Mice containing vibration feedback can also be used, as the haptic system is very responsive to vibrations, however it, as with all the above devices, suffer from the inherent problem of using a single point contact device: information about only a single contact point is *extremely* poor in conveying information about an object [33]. Methods to provide directional cues have been tried (e.g., [32, 20]) but their success is unclear.

Another possible improvement on a point contact interface is to use multiple point contacts, as it potentially provides significantly more information in parallel: Lederman and Klatzky [33, 34] found that, at least for 3-D object recognition, the use of multiple fingers of a hand was an important source of information. This is very difficult to do with sonification, although possible with force-feedback and vibration feedback. However, no cost effective devices have been proposed yet. Multiple finger force-feedback gloves and vibrotactile feedback gloves have been developed for virtual reality

(e.g., [35] and Immersion Corporation) but are prohibitively expensive for this application.

2.2.2.2: Distributed Displays

Lederman and Klatzky [37] found that, when using a single finger, eliminating spatially distributed information to a person's fingerpad reduced to chance their ability to determine even coarse 2-D bar orientation with that one finger. This suggests the importance of spatially distributed displays within a finger for portraying haptic graphics. Many types of devices are being investigated for these displays (for recent reviews see [38, 39]); however, most of these are still in the development phase and have various hurdles to overcome to become practical. In terms of commercial devices, large, refreshable displays are available but are prohibitively expensive and, potentially, difficult to maintain.

To overcome the size problem, some smaller displays have been combined with an optical or position sensor acting as a scanning device over the visual image (e.g., the Optacon; the VT Player by VirTouch). Although the Optacon has a loyal following, the device is both expensive and difficult to learn how to use. This latter is, to a high degree, due to the fact that the tactile information is displayed on a square grid of pins spaced 2mm apart and vibrating at a frequency of 230Hz. Unfortunately the primary tactile receptors that this vibratory frequency excites have poor spatial resolution, significantly reducing the effect of using multiple pins. In addition, the kinesthetic and tactile information go to two separate hands, which increases the difficulty of an already demanding task. The commercially available VT Player, in contrast, is relatively

inexpensive and excites the most appropriate receptors (i.e., the Slowly Adapting Units). However, one difficulty with this device is the relatively slow update rate, which creates a time lag between the kinesthetic and tactile information.

Two such aspects are the lack of spatial or temporal collocation between the kinesthetic information and the tactile information. Lack of temporal collocation is expected to create the most difficulty as a noticeable temporal delay (as exists in the commercially available VT Player) creates a different spatial inconsistency between the tactile and kinesthetic information depending on the direction and speed of motion. In contrast, devices lacking spatial collocation, such as the Optacon (where the kinesthetic information is provided to one hand and the tactile information to the other), are at least consistent in their spatial mismatch. However, the transformation needed to align these two types of information still increases the cognitive load on an already taxed system.

Although, a single point contact device has shown to be clearly insufficient, it has been suggested [33] that multiple point feedback devices to a hand may be able to make up for the lack of distributed information on a finger. This is particularly expected to be true if solid, textured surfaces are used, as the fingers are particularly effective in process material properties in parallel [36]. The device was made to facilitate the use of multiple, independent devices simultaneously to provide multiple point feedback; thus, maximizing the amount of parallel information obtained within design constraints.

3. Device Design

The device purposed is designed to be a solution to the problem with rendering 2-D graphics for individuals who are blind or otherwise visually impaired. It was designed to not only with consideration of the haptic system, but with consideration of the blind and deaf-blind community as well. The devices developed all represent a hardware solution to the problem, since software solutions necessitate the user having a computer; and although the last prototype discussed currently requires a computer, it could be modified to work without one. The hardware solutions proposed enable the device to have dynamic display capabilities through active sensing of a graphic by the optical sensor and associative control of the haptic actuator. The increased versatility of a dynamic display permits the device to be used on any number of graphics without costly and time-inefficient static image production. This versatility also translates to increased cost-savings, as it reduces the immediate need for tactile printers, swell/puff paper, and other materials associated with static tactile image production.

Another device design choice universal to all the prototype models is the choice of vibrotactile actuation as the means for providing haptic feedback. Vibrotactile output was initially chosen for its cost-efficiency, as small pager-motors can be readily obtained for around $1US. Even the speakers (either voice-coil or piezoelectric speakers) later chosen were far more cost-efficient than using pin-arrays or force-feedback mechanisms. Cost-efficiency is defined in this case as being the ratio of the expected actuator performance (in terms of user perception) to the cost of implementing the actuator; in assessing the

expected actuator performance, only informal observations were used, rather than build separate systems and quantitatively compare them. This was seen as an unnecessary task, as ultimately the device will be benchmarked against raised-line images, the standard for tactile graphics.

Vibrotactile feedback is also an advantageous choice in that vibration sensing is crucial in the perception of fine roughness in material surface texture [3, 45, 46]. By providing different vibrotactile outputs comprised of different vibration parameters, different textures can be simulated. Also, texture is one of the tactile perceptional dimensions which benefits from early-stage parallel processing across multiple fingers [36]. Additionally, texture has remarkable discrimination [2, 36, 46]; however, care must be taken to make sure the vibrotactile stimuli are appropriately discernable [50, 52, 53]. Through the use of temporally encoded textures, the feedback of device possesses both increased perceptive bandwidth through the use of multiple fingers versus contour-following and high accuracy in stimulus discrimination.

These considerations and choices formed the basis for all the designs, though many of the design points developed concurrently with the prototypes. Initially, when vibrotactile actuation was chosen as the haptic feedback mechanism, little consideration was given to the fact that textures could be processed in parallel. However, after the stylus was designed, it was determined that switching to a glove-based design would enable the use of multiple fingers and the increased perceptive bandwidth associated with parallel processing. Thus, the purpose of the second prototype was to allow for multiple fingers;

however, during its development it was determined that the pager-motor actuators were too limited in terms of their ability to generate different outputs. Also, the optical sensor had its limitations, as it was determined that through the use of colors that the device could be used more efficiently by segments of the population with some residual vision. Therefore, the third prototype focused on implementing a new color sensing optical element along with a piezoelectric actuator capable of output a wide range of textures.

3.1 Design Considerations

Important design characteristics of any system for the haptic perception of 2-D graphics are: 1) the spatial resolution that the device is capable of rendering (for perceiving 2-D geometric information as opposed to texture); 2) the temporal concordance of the tactile and kinesthetic information (i.e., any signal latency incurred by the device); and 3) the spatial concordance of the tactile and kinesthetic information. The spatial resolution should match or be as close to human perception as possible, so as to relay geometric information as accurately as can be, but is not an absolute necessity. The values given in the literature about human perceptual thresholds are mean threshold values, with the total population having a distribution of spatial resolution thresholds—dogmatically trying to obtain a specific spatial resolution is not likely to improve device performance. For temporal concordance, the signal latency has to be sufficient to not cause any noticeable spatial inconsistencies in conditions of typical scanning velocities, which range from 20 to 80mm/s. This means the acceptable limit of latency is dependent on the devices spatial resolution and the maximum scanning velocity. If a maximum scanning velocity of 100mm/s is assumed with a spatial resolution of 1mm, then the

maximum allowable latency is 10ms. Spatial concordance between the tactile feedback and kinesthetic information is also important; the spatial difference between the point of contact for the tactile feedback and the optical scanner should be as minimal as possible. This is of particular importance if the device is to be expanded to multiple point contacts; otherwise, the mental effort involved in correlating the feedback can become very cognitively taxing.

Other important criteria are that the device has to be intuitive, affordable, portable, safe and comfortable to use. An individual's intuition towards a device is a characteristic of the usability of the device, specifically, how easily learned are the device's functions. This can be measured by the length of time it takes to train individuals to use the device and the ability of individuals to successfully use the device. To limit the time necessary to train individuals, the device should have a low cognitive demand. This can be achieved through the simplicity of its operation and the ease of the perception of its output or interface. The cost of the device is an important criterion as it is one of the key components for acceptance by individuals who are visually impaired. As of 2002, only slightly more than half of the individuals who were blind or visually impaired had employment, with a medium income of just less than $16,000US annually [13]. Therefore, the target manufacturing cost for the device was set to be no more than $100US, so that it would be an accessible option for all individuals who are visually impaired. Another criterion for the device is that it has to be portable, as portability contributes greatly to the device's overall usability. The device's ability to be easily carried from home to the work environment or school is a crucial factor to contributing to

the independence of the user, a major issue with individuals who are blind or visually impaired. It also contributes to the accessibility of the device: by being lightweight, it helps children adopt it as a tool they can easily use and carry with them. An added bonus would be if the device is, or could be made stand-alone, requiring no other equipment, such as a computer or smartphone, to work. Even as ubiquitous as computers are in this day and age, it is difficult to always have one in every situation where visual images are encountered.

The final criteria for the device are that it be safe and comfortable to use. The principle safety concern for most of these devices is eliminating the possibility of the user developing HAVS (Hand-arm vibration syndrome) from using the device for prolonged periods of time. HAVS is brought about by time-dependant, frequency-dependant exposure to vibrations. The stronger the vibration amplitude, the shorter the safe time of exposure before HAVS is a problem. Certain frequencies also present an increased risk for triggering HAVS [8]. However, current research has shown a need to revise these standards, as the frequency weighting proposed does not match experimental data [9]. This will be an important component of selecting the "textures" to be used, not only with regard with their amplitude, but their frequency as well. The comfort of the device is another important consideration, as individuals are less likely to adopt a device that is not comfortable to use. Unfortunately, this design criterion is harder to accomplish when making initial prototypes, as the additional time needed to make rounded edges and polished surfaces can greatly increase the cost of manufacturing. Essentially, as long as

the prototype device is judged to be bearable to wear for testing, then it is accepted with further comfort related design issues being addressed at a later point.

3.2: First prototype (Prototype 1.0: Stylus Design)

The goal of the first prototype was to develop a device capable of detecting borders or outline figures and to serve as a proof-of-concept that the design theory worked. The detection of borders or outlines only requires binary feedback, as border detection is a Boolean operation (is there a border: true or false). This simplified the device operation, since it only needed to detect highly contrasting differences (white versus black) in the graphics used. Thus, the first design was kept simple, so as to be cost-effective, and easy to use. The two principle components of the prototype were: (1) an optical sensing circuit, which was used to parse widely distributed levels of grayscale contrast (which could either be on paper or a computer screen) into either a "white" background or "black" lines or solid shapes, and (2) a mechanically vibrating actuator that stimulated the same area of the hand used in sensing when the contrast detected was "black". A push-button switch acted as an interrupt between the battery and the rest of the circuitry, acting as an on/off switch for the device. The components were encased in a pen-like form and the device operated similar to a pen: the user simply had to press the device against the surface and trace over the graphic.

3.2.1: Choice of System Components

To keep the kinesthetic information collocated with the vibratory feedback, several different types of sensors were considered for sensing the point of interaction on the graphic. Several position-sensing systems were examined, such as using an absolute-

position detecting RF transmitter/receiver pad, similar to what is used in graphics tablets such as the Wacom® Intuos® 3, or through the use of dual-axis accelerometers. However, a simpler, more temporally responsive (i.e., in contrast to a Wacom tablet which has a temporal response of only about 20Hz) and more cost effective alternative was chosen: a simple optical analog detector that could determine the contrast of a visual graphic made of raised lines and/or solid surfaces placed on a computer screen or paper. A photo-interrupter was chosen over other simple optical circuits, as the photo-interrupter contains both a light-emitting diode and a phototransistor in a discrete package, enabling the device to work without an additional lighting source; unlike photo-resistors or phototransistors, which require a photon source such as a light-emitting diode in order to generate a signal.

In order to work with an optical sensor, the visual graphic needs to be presented in a simplified form for translation to haptics (i.e., with any encoding of information, magnification and/or simplification already performed). However, this can be easily done by a computer, and allows individuals with low vision to use any residual visual as well to interpret the graphic. The visual presentation also allows a sighted observer to see the same information presented dynamically to the user. This is expected to allow individuals who are blind or visually impaired to be able to communicate more effectively to someone who is sighted about the information. Using an optical sensor also means that the workspace is restricted by the size of the visual graphic, not the tactile device: an added advantage when examining large paper graphics, such as floor plans.

Vibrotactile actuation was chosen as the method for producing haptic feedback due to its low cost, high performance ratio and the ease of its implementation (i.e. it did not require a microcontroller, timing circuitry, or software control) relative to the rest of the device circuitry. A pager motor was selected as the specific vibrotactile actuator, as it is the cheapest available vibrotactile actuator, it gives strong vibration feedback, and works relatively well using a binary output control.

The overall device was housed in a stylus design; although a glove type mimics natural haptic exploration, the stylus device was far easier to implement as a proof of concept device, while maintaining concordance. A mouse-like device was rejected, as mice often convey a poor sense of spatial location to the user, as the point contact becomes the entire hand rather than a single finger with a stylus. Figure 2 shows the intended stylus design for the first prototype.

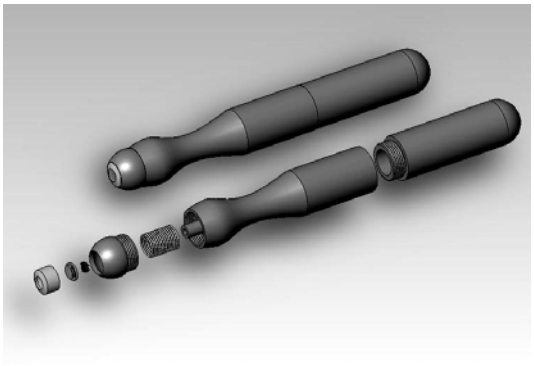

Figure 2, the first Prototype Design

In this figure both the exploded view and the assembled view of the case prototype are shown. The tip consists of two main pieces, one metallic colored the other blue, with the lens and the photo-interrupter following the metallic colored piece, then the

secondary blue piece that screws onto the next section, and the spring. The middle piece contains both the motor in the slimmer section, and the circuit board in the wider section. Finally, the back piece contains the battery; however, it would probably be easier to have the back section open up closer to the end to make it easier to run wires between the battery and the circuit board. However, it was decided that a glove design had more promise for future development of the device, so the stylus design was abandoned after the first model.

3.2.2: System Design and Method of Operation

The photo-interrupter consisted of an infrared photo-diode adjacent to a correspondingly sensitive photo-transistor in a discrete package. It operated by emitting infrared light from the diode, which hits a nearby surface (the photo-interrupter used has an optimal 0.6-0.8mm range, but works up to around 0.4-1.4mm) and is either absorbed or reflected by the surface, depending upon surface characteristics, of which saturation is of particular importance. The light that is reflected from the surface back onto the phototransistor provided a voltage potential at the transistor's base by the photoelectric effect. If the amount of reflected light reached a certain potential, then the transistor will turn on, allowing for current to flow across the collector-emitter bridge. As the potential at the transistor base increased, the amount of current allowed through the collector-emitter increased, corresponding to a drop in the collector voltage. Thus, the drop in collector voltage produced an analog signal of changing voltage potentials that correspond to changes in reflected light hitting the transistor base.

Initially, the circuitry was designed to be used with op-amps to first buffer the signal, provide a voltage gain, then remove the DC offset through a summing op amp. The first "working" circuit design is shown in Figure 3.

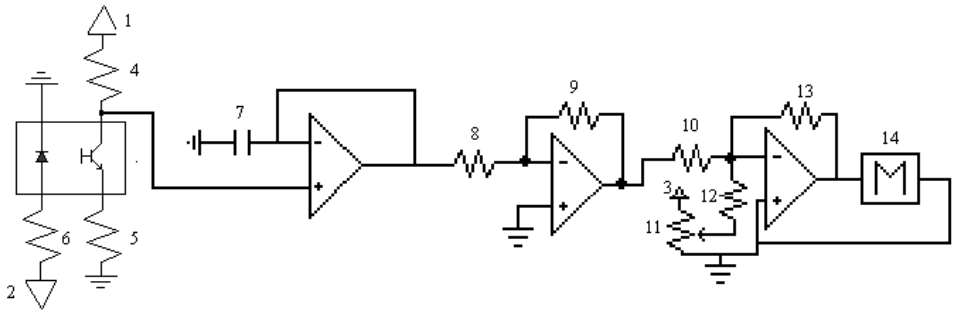

Figure 3, the initial circuit for the first prototype

The drawback to this design was that the motor would draw enough current to fry the AD744 op amps. Obtaining op-amps that could run off of low supply voltages, while allowing for the current output necessary to drive the motor, was not as cost-effective of a solution as running the whole circuitry through diodes and transistors. In the new design, the analog signal from the photo-interrupter outputs through four NP-diodes to remove its baseline voltage. Then the signal is sent, to a Schmidt trigger circuit. The Schmidt trigger was used to improve the spatial resolution of the device, turning on the vibrating actuator more effectively at finer resolution. This was actually necessary, because without it, the voltage output from the photo-interrupter was insufficient to be used with a power transistor pair to drive the motor. The motor would turn on inconsistently and at much lower rotational speeds than it would for thicker lines or patches of solid black. The hysteresis voltage for the trigger was set to decrease the temporal response by having the

second trigger voltage only slightly less than the first. The Schmidt trigger then fed into a modified Darlington-pair to boost the current allowed to the motor. The final circuit is shown in Figure 4.

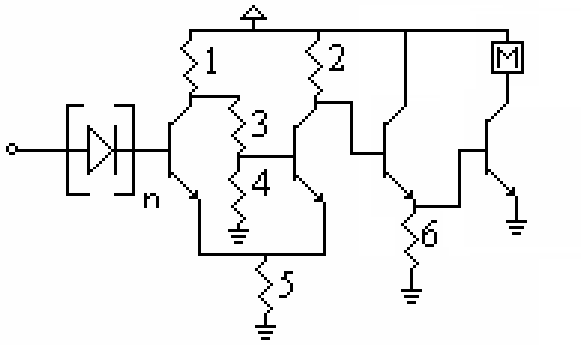

Figure 4, the final circuit design for the first prototype

A single pager motor was chosen for the vibrating actuator for its very good temporal response, size and affordability. A counterweight was attached to the pager motor, which vibrated the whole actuator system at approximately 203 Hz. Although this meant that primarily tactile receptors with poor spatial resolution respond to the vibration information, their resolution is sufficient to resolve one actuator per finger (as is proposed). Note that any spatial resolution less than this would be wasted, as with the Optacon, due to the inability of the tactile perceptual system to determine spatial detail at this high of vibration frequency.

An additional mechanism was necessary to detect when the device was in contact with the graphic: otherwise, the photo-interrupter would provide a high signal when not in contact with the medium (as if it detected a dark contrast area) turning the motor on erroneously. Since a user-controlled on/off mechanism was already needed on the device,

the idea was to activate the device with a contact sensor, so that it would only turn on when being used. The most cost-effective option available was to use a push-button switch, which had the photo-interrupter and lens mounted on its contact surface. This setup causes power to the device to be interrupted anytime the device is lifted off the workspace surface.

Finally, a 6V RadioShack Lithium battery (2CR-1/2N) was used to drive the circuits, allowing for 4.7 hours of operation in the on position (with the motor drawing 250 mA). However, the life-time of the device is expected to be much greater than this due to the little amount of current drawn in the off position and the typically short intervals of motor operation during graphic exploration.

3.2.3: Discussion

Using largely off-the-shelf components from hobby shops and electronic wholesalers reduced the cost significantly. Further, using the device required few steps, reducing the learning curve for the device. The design was made kinesthetically and temporally concordant by: (1) providing the appropriate tactile feedback, based on the graphic, to not only the same hand, but the same part of the hand, as the kinesthetic information being obtained, and (2) using vibratory actuators, such as pager motors or audio transducers, which have a very short temporal response and are also cost-effective and relatively small. Overall, the prototype achieved its goal of being a proof-of-concept for an affordable and highly portable device capable of detecting borders and outlines in simplified 2-D graphics.

3.3: Second prototype (Prototype 1.5: Glove Design)

The goal for the second prototype was to shift the design from a stylus model to a glove model, to allow for expansion to 5-finger device use and to mimic a more natural form of haptic exploration; this is intended to allow for parallel processing of vibrotactile stimuli across multiple-fingers. The photo-interrupter, the circuitry, and the motor remained the same for this prototype; only the device casing changed, so it is referred to as version 1.5. Unlike a stylus design, a glove design allows for multiple interaction points. A worn device has many advantages: (1) by placing the sensing elements on the fingertips, it mimics natural haptic exploration; (2) the vibratory motors can be mounted on the glove near the sensing element for concordance between optical sensation and tactile stimulation; and (3) it allows for independent device movement of multiple devices without requiring the user to have to hold them. The use of an optical detector and vibratory actuator are also easily extended to multiple simultaneous but independent devices placed on a person's hand, which was an important requirement for our design. Figure 5 shows the glove-styled prototype: the circuitry and battery for the device are housed in a small project box from RadioShack (A), wires output to the pager motor (B), and input from the case housing the optical sensor and the push-button switch (C). Aside from the custom manufactured case for the optical sensor, all the components cost less ($12) than a pair of the work gloves they attach to ($16).

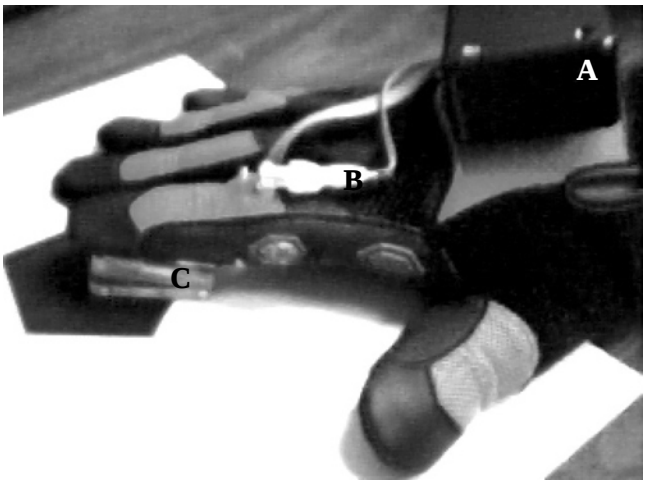

Figure 5 the glove prototype model

In addition, the nature of the lens for the optical sensor was changed to incorporate the push-button switch and to increase the stability of the lens-sensor interface. In the previous prototype, the stresses placed on the lens during exploration would sometimes damage the photo-interrupter leads or the contact with the push-button switch. To eliminate this, a larger case was decided by then undergraduate assistant Justin Owen to house the optical sensor and push-button switch that would reduce these stresses and that could be mounted onto a glove. This design is shown in Figure 6. The advantages of the design are that it allows for easier access to the photo-interrupter and push-button switch for modification or repair. Also, since the case was made out of a transparent acrylic plastic, it did not require a separate lens.

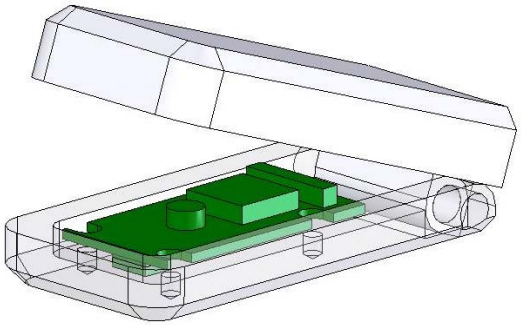

Figure 6 the sensor case design model

3.3.1: Spatial Resolution

To adequately evaluate the spatial resolution of the device, a series of images were designed to test the principle aspects of its spatial behavior. The two key components examined were: (1) the thinnest line for which the device would trigger and (2) the smallest separation between two lines that the device could distinguish. Since it was known that the sensor viewing field was not symmetric with respect to the horizontal and vertical axes, it was necessary to test how the orientation of the sensor relative to the image affected these values as well.

The first test of the device was to confirm that the Schmidt trigger functioned as expected, by first using horizontally-oriented scans with lines varying in thickness. The lines were made long enough (>10cm) so that the sensor portion of the device always passed over them. The line thicknesses tested were 0.33mm, 0.50mm, 0.66mm, 0.75mm, 1.00mm, 1.33mm, 1.50mm, and 1.66mm. Each line was independently measured and printed at 1200 dpi on a laser printer to insure uniformity and accuracy. The device was passed over each line 50 times at approximately the same speed of 50 mm/s. To evaluate the success of a test, visual and haptic inspection of the motor were used to see if it

triggered. The motor never triggered for lines of 0.33mm or 0.50mm thickness (0 times out of 100 total tests), while lines of 0.66mm, 0.75mm, 1.00mm, 1.33mm, or 1.50mm in thickness trigger nearly 100% of the time (250 times out of 250 total tests). Reorienting the lines in a vertical position yielded the same results; the motor never triggered for lines of 0.33mm or 0.50mm thickness (0 times out of 100 total tests) and 100% of the time for lines of 0.66mm, 0.75mm, 1.00mm, 1.33mm, or 1.50mm in thickness (250 times out of 250 total tests). This test showed that, for single line detection, the device is insensitive to orientation within the plane of exploration, with sensitivity to line thickness of about 0.66mm.

The second test of the device was to establish the smallest separation between lines that the device could detect, which defines the spatial resolution of the device. Again, a series of lines were made, but this time with varying separation distances. Each set of lines contained three lines of equal thickness with separation distances of 0.5mm, 0.66mm, 1mm, 1.3mm, 1.5mm, 1.66mm, 2mm, 2.5mm, 3mm, 3.5mm, 4mm, 4.5mm, or 5mm as measured from the line edge to the next line edge. Each line separation distance was repeated for line thicknesses of 0.5mm, 0.66mm, and 1mm. The purpose of this latter manipulation was to see if the thickness of the line, as well as the separation between it and the next line, affected the results. The device was tested using both horizontal and vertical scans, as it was expected that the increased vertical field dimension would increase the necessary separation. The device was passed over each set of lines 20 times. Both the change in the base voltage on the motor's current amplifier (as monitored through an oscilloscope) and physical inspection of the motor were used as methods to

evaluate successful triggering of the motor in an on-off-on pattern for this test. The minimal separation needed for horizontal scanning was found to be 1.33mm for 0.5 mm (14 times out of 20 tries) and 0.66mm thick lines (20 times out of 20 tries), and 1.5mm for 1mm thick lines (20 times out of 20 tries). This suggests that there is somewhat of an interaction between line thickness and spatial resolution in the horizontal direction. In contrast, the spatial resolution in the vertical direction was found to be invariant: the minimal separation needed for vertical scanning was found to be 2mm for any thickness line (60 times out of 60 tries). Spacing less than the minimal separation caused the motor to operate continuously, as if the separate lines were a solid rectangle. It is suspected that the differences between the horizontal and vertical scanning characteristics are due to the rectangular shape of the phototransistor base, which is longer in length (the axis that lies in the sagittal plane with respect to the rest of the device) than in width (the axis that lies in the frontal plane with respect to the rest of the device).

Thus, the overall functional spatial resolution was determined to be a thickness of 0.66mm for single lines and an absolute minimum line separation of 2mm, which should be used in designing the graphics used to minimize user confusion due to false sensor readings. It should be noted that the center to center line separation of 2.6mm from these specifications is not too far from the spatial resolution of the tactile system which was previously stated to be approximately 1mm. [See Appendix 1, Table A for data]

3.3.2: Temporal Collocation

The time delay between the optical sensor detection and the motor triggering was measured, as it determines the temporal collocation of the device, an important factor

when assessing the device's ability to quickly transition from one mode of operation to another. This prototype only has two modes of operation, on and off, so the time necessary for the motor to turn on and the time necessary for it to turn off were the two cases examined. The time delay was measured directly using an oscilloscope to observe the output of the photo-interrupter on one channel, and the output of a small MEMS accelerometer glued directly onto the motor casing on another channel. The motor casing was attached on the glove so that the shaft was allowed to rotate freely and uninhibited by any connecting wires or the accelerometer. The device was then scanned over a solid black rectangle of sufficient size to trigger the motor, and the time delay was then measured by comparing the time difference between either peak voltages (going from off to on) or baseline voltages (going from on to off), for both channels using the oscilloscope. Figure 7 below shows the oscilloscope screen during a measurement of the device triggering from an off-state to an on-state: the top waveform (Channel 2) is the output of the accelerometer attached to the motor and the bottom waveform (Channel 1) is the output of the photo-interrupter.

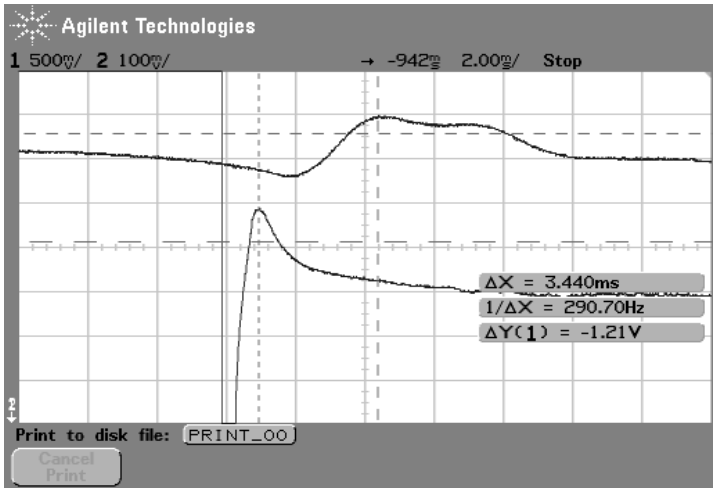

Figure 7 the calculation of temporal delay on the oscilloscope

The average time delay for ten trials for each change (twenty total) was 3.4 milliseconds, with a variance of less than 1 millisecond: well below generating any noticeable spatial inconsistencies for typical scanning velocities of 20-80 mm/sec. Both the optical sensor and the accelerometer have negligible delays associated with them, with the larger of the two delays, the accelerometer, being on the order of a tenth of a millisecond. [See Appendix 1, Table B for data]

3.3.3: Preliminary testing with 2-D Graphics

Preliminary testing used basic geometric shapes, simple mathematical curves and x-y graphs to examine whether the usability of the device with one finger was acceptable, before extending it to multiple fingers. The purpose of this test was to have a basis for judging the fundamental usability of the device as a proof of concept. Prior to this, no formal testing had been performed using the device, only very informal tests using the device on graphics presented on normal paper and duplicates for comparison with raised line drawing techniques on swell paper. The primary concern with these informal tests was to qualitatively see how well the device worked and make any necessary improvements. These tests showed a need to reduce the force necessary to close the push-button switch by changing the push-button switch to a different model. They also showed a need to replace the lens structure with a different case design.

For the more formal tests, six sighted subjects within the Department of Biomedical Engineering, between the ages of 20 and 25, volunteered for the test. Two of the subjects had some experience with the device, while the remaining four had no prior experience.

Each of the six subjects were taught how the device works and allowed to freely explore a 15 cm x 15 cm solid square on a 21.5 cm by 28 cm sheet of paper. The subjects were then blindfolded and told to name the shape in three different images, each presented in an individual trial. The images were randomly selected from a pool of images which included a solid triangle and a solid circle; an outlined square, triangle, and circle; and a parabola and a third-order polynomial curve. The dimensions for each of the outlined and solid images were approximately 15 cm by 15 cm, whereas each of the curves varied in overall length. The sheets of paper they were on were 21.5cm by 28cm and line thicknesses for all non-solid images were 5mm. Subjects were not informed of what images were in the pool ahead of time. For the latter group of images, determining the overall shape of the curve substituted as a correct answer in place of naming the shape. Although subjects were not given time restrictions during their exploration of an image, no subject took more than three minutes for any single image in the task.

The subjects were then asked to explore one of two x-y graphs (both the full length of a 21.cm by 28cm sheet of paper with line thicknesses of 5mm), and answer questions that required them to point out the absolute and local maxima, minima, and compare slope magnitude for various graph features. The subjects were first allowed some time to explore the graph, before they were asked to determine the aforementioned characteristics of the graph, one-by-one. Subjects were not timed for this exercise and typically people took anywhere from three to six minutes to complete the task. After the test, subjects were asked to record any comments that they had about the device.

The results in Table 1 show that, for the preliminary tests, the subjects had a slightly higher than 50% accuracy (59.2%). Although, no quantitative comparison evaluation was made to existing methods such as raised-line images or other existing devices, qualitatively the results seem promising for the device's extension to multiple fingers.

Table 1:

Image	N	Mean % Accuracy
Solid Triangle	3	67
Solid Circle	3	33
Outline Square	1	100
Outline Triangle	2	100
Outline Circle	3	33
Parabola	3	67
4-Point Spline	3	33
Graph 1		
Maxima	3	33
Minima	3	67
Slope	3	0.0
Graph 2		
Maxima	3	33
Minima	3	100
Slope	3	67

As a side comment, noting the results of Thompson and her colleagues [2, 7] that suggested the higher effectiveness of solid images over outline drawings, the results from

using raised outline drawings to solid images was compared. A matched-pair t-test was performed (df = 5) to compare the two types of images, yielding a p-value of 0.613, although the power was low due to the limited number of subjects for any one test. Since at this time work on the third prototype was already underway, it was decided that conducting more thorough testing on this prototype was an inefficient allocation of time.

3.3.4: User Safety

In terms of safety, the prototype was a low power electronic device with the active components encased in a plastic enclosure on the back of the hand. The vibration feedback was provided by small, commercially available pager motors. However, as the pager motors were being used in a different setting than the one for which they were designed, there was a potential risk of hand-arm vibration syndrome (HAVS). Early symptoms include: chronic tingling and numbness in the fingers, not being able to feel things properly, loss of strength in the hands, and the fingers going white (or paler). Symptoms typically appear after a few months to a few years of exposure.

The safety of the device was evaluated, before using it in the preliminary testing (above), in terms of the International ISO 5239-1 Standard [8] for determining the safe exposure to vibrations in the work place. The acceleration of the motor was recorded using an Analog Devices ADXL203 Dual Axis Accelerometer, which was mounted on the lateral side of the glove near the interphalangeal joint, approximately 1cm away from the motor. The accelerometer was oriented so that the two axes lay in the plane of the motor's rotation, perpendicular to the motor's shaft. Formal data was not collected for the third axis, which would lie parallel to the motor's shaft, as acceleration along this axis

was informally found to be insignificant. The two channel output from the accelerometer was input into LabView, where both channels were filtered with band-limiting and frequency weighting filters, in accordance with the International Standard ISO 5349-1. Four trials of 10 samples each were recorded.

The average magnitude of acceleration of the vibratory motor was measured at 2.411 m/s^2, with the maximum recorded acceleration at 2.420 m/s^2. The single-axis average accelerations were 1.625 m/s^2 and 1.781 m/s^2. Figure 8 is the pre-filtered Fast Fourier Transform obtained from LabView showing that the motor vibrates at a fundamental frequency of approximately 203 Hz with several harmonic frequencies; the dc offset shown in the picture should be ignored, as it is a component added by the accelerometer's zero-g base voltage. With all of the acceleration frequencies well beyond the 15.915 Hz frequency known to exacerbate HAVS, and the total RMS acceleration magnitude below the threshold for precaution at any exposure duration (including for the entire work day), it was surmised that the motor would present minimal to no risk for the user.

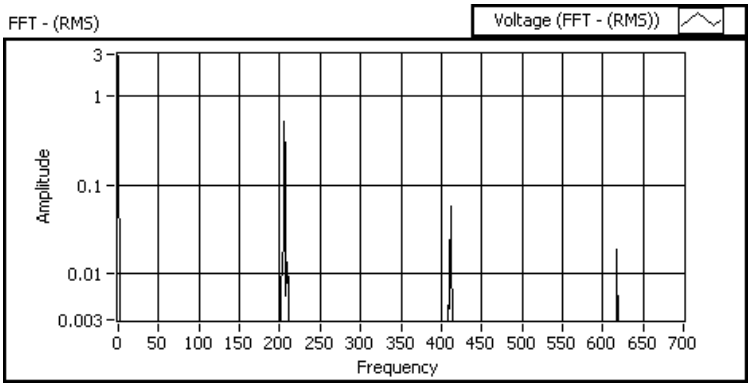

Figure 8 the FFT of the accelerometer input

At a value of 4 hours of total daily exposure (a value which greatly exceeds even the most liberal estimate of use) the daily exposure value of our device is 1.697 m/s^2, using the ISO 5349-1 guideline for calculating magnitude. This is well below the 2.5m/s^2 acceleration level at which action needs to be taken.

However, during the preliminary testing on actual graphics, 3 out of 7 of the subjects experienced a slight discomfort when using the device. They described the feeling as "a tingling sensation" in the finger and hand to "numbness" in the finger. The numbness felt was further described as a loss of tactile sensation, excluding static pressure. Subjects who experienced the discomfort said the onset came shortly after exploring the large, solid images. Some subjects further went on to say that exploring the outline images gave them no discomfort but the longer percent use of the device when scanning the solid images did. None of the subjects experienced the discomfort long after ceasing to use the device (5 minutes or less). It is not clear whether the discomfort resulted from adaptation to the vibratory stimulus, or from an associated effect of the vibration-related disorder, similar to an acute phase of HAVS. While the acceleration test data indicated that the vibration magnitude was insufficient to cause long-term problems such as HAVS (using the ISO 5349-1 Standard), the Standard does not rule out the possibility of causing short term discomfort for some users. Some insensitivity to high frequency vibration (i.e., adaptation) was noted with the discomfort, but it was not easily apparent whether it was the main cause of the discomfort, or a side-effect of a more significant problem.

Also, while the International Standards ISO 5239-1 and ISO 5349-1 provide a guide for the maximum daily vibration exposure for an individual by proposing a frequency-weighted algorithm for measuring the acceleration magnitude of these vibrations, our tests have provided anecdotal evidence that they may not sufficiency address the problem. This hypothesis is supported by work of Dong and his colleagues [9] who found that the energy absorption distribution for vibration is dependent on the frequency, with higher frequency (>100Hz) energy being absorbed more locally. Further, they suggest that for frequencies higher than 16Hz, with some exceptions, "the relative weighting is higher than that of the ISO" [9]. This would explain the findings, but more testing would be needed to prove it to be the case. Regardless of whether the discomfort resulted from adaptation to the vibratory stimulus, or from a HAVS-like effect resulting from the vibration amplitude, discomfort associated with tactile stimulation was not acceptable for the device. The solution to this problem was to reexamine the actuator choices, which lead to the new of choice of using piezoelectric buzzers. The buzzers have an advantage over the solenoid motors in that they produce lower amplitude vibratory stimulation; however, this requires them to be placed directly against the skin to be felt, unlike the solenoid motors. Fortunately, the thin casing for piezoelectric buzzers facilitates such placement, which improves the kinesthetic coupling of the optical sensor-actuator system. This modification to the design was implemented in the third prototype model.

3.3.5: Adaptation

In addition to HAVS, the user's tactile perception is expected to become desensitized over time [10], potentially making use of the device frustrating. In one test, a subject

used the device for a period of two hours, with the input to the motor bypassed with a constant voltage so that the motor would stay on. Throughout the two hours, the subject showed no adaptation or discomfort to the vibratory stimulus. However, during a second test, 3 of the subjects mentioned some discomfort, with loss of sensation being one of the mentioned effects. Incidentally, the subject who tested for adaptation also participated in the previous study to identify objects and did not report feeling any HAVS-like symptoms. Currently, it is not clear whether this resulted from receptor desensitization or from an effect similar to, if not, HAVS.

3.3.6: Discussion

This prototype design accomplished its primary goal of transferring the stylus-based design to a glove-based design; however, the device was never actually expanded to use multiple fingers. The reason for this, as previously stated, was because expanding the device would have involved extra work that ultimately would have been in vain. It was determined that the pager motor was an inadequate actuator choice for providing multiple vibrotactile stimuli and it had the added disadvantage of producing noticeable discomfort among some of the test subjects. The optical sensor also presented issues as the function of the device was expanded to color detection, because the photo-interrupter could only detect changes in saturation, but not hue. Nevertheless, the overall design was success in terms of having adequate spatial resolution, temporal collocation, and kinesthetic concordance.

3.4: Third prototype (Prototype 2.0: New Glove Design)

The main switch in the third prototype was to the ability to controllably generate arbitrary temporal waveforms on the device display. This allows the ability to generate texture-like patterns. The advantage of using textures in displaying graphics is that, with a multi-finger display, they can potentially be processed, in parallel, across multiple fingers at a time: thus, speeding up the exploration process of the graphic. The idea is that different regions of the graphic would have different textures, allowing an individual to obtain coarse information about 2-D form by placing all five fingers on the graphic, rather than the laborious serially processing of geometric information. Further using texture to encode information about the separation of an object into parts and the orientation of each part has a significant additional advantage as these two aspects of raised line drawings are the most confusing to interpret in the graphic [55]; the advantage of such a method has been shown in manual relief pictures by Thompson and her colleagues [7].

With the switch to using textures, if optical sensing is to be used to interpret a graphic, then some means of locally encoding that information in the graphic, interpretable by an optical sensor, is needed. (As opposed to position sensing where the conversion can be done algorithmically based on position. However, with position sensing, where the finger is in space is not directly collocated with the location on the graphic, which makes the use of any residual vision less easily interpretable. In addition, using position sensing is less easily transferable to multiple fingers.) Different grey scale values could be used to represent each texture visually. However, the use of color is

more advantageous than grey scale as high contrasting colors, which are also good for detection by a physical color sensor, aid in the detection of form by individuals with low vision.

The new design does not differ from the previous version in its overall design operation: an optical sensor is used to detect light reflecting off a surface and this is translated into vibratory feedback to the finger. However, the "glove" design has been replaced with a small plastic case worn under the finger and held in place by an elastic band that wraps around the finger. Currently, the system utilizes a Labview program to control the device; thus, reducing the need for much of the circuitry. This eliminated the case worn on the back of the hand in the previous design and removed the need for a glove to mount parts. However, ultimately the design will shift to use an embedded controller, or other hardware, versus a whole computer, at which point the case on the back of the hand will most likely return to house the additional hardware. For the optical sensor, the photo-interrupter has been replaced with a Red-Green-Blue photodiode sensor. Both it and the haptic actuator are now housed inside the case along with the circuitry to filter and amplify the signals. Cables stemming from the case provide three channels for analog input (corresponding to the three channels of sensor output) and one channel for analog output (corresponding to the signal that drives the haptic actuator) to a National Instruments DAQ card. Additionally, three power cables also come from the back, to provide the op amps with ±12V supply voltage and a ground connection.

3.4.1: Device Components

3.4.1.1: The New Optical Sensor

One of the most important differences between this device and the previous iteration is the added ability to sense color. This ability serves two purposes: 1) it allows for added usability for people with low vision, for whom the haptic feedback supplements the visual image they see (through the use of highly contrasting colors, as recommended by the guidelines for producing images for individuals with low vision [12]) and 2) it allows an easy means to encode information about the separation of an object into parts and part orientation in the actual visual image, which can then be translated by the device into textural information. By programming the device to encode a unique texture for every color it is set up to detect, the system is set up to easily render graphics haptically with the use of textures. However, in order to effectively use color with the device, a new optical sensor was needed.

As most optical sensors require a light source to operate, it was decided for this prototype that the device would only work with a computer monitor, with the monitor being used as a light source. Working with an LCD monitor, in particular, facilitates things even more, since the flat surface is easier to explore and the monitor can be easily oriented on its backside, a more comfortable position for the hand. In the future, a light source can be added to the device to enable it to work with paper with encoded color diagrams as well.

A variety of optical sensors were analyzed for their potential use. Most of the sensors examined were photodiodes of some type. Specifically, the BCS series by TDK and the

OPT101 series by Texas Instruments photodiodes, as well as the Agilent HDJD-S831-QT333 CMOS chip, the Hamamatsu S9706 Digital RGB Color sensor, and the Hamamatsu S9032-02 Analog Color sensor. The BCS series had two disadvantages: they were sensitive to only one color (blue), requiring an inverse transform to normalize the response, and had too large of a receptive field (~4mm). The OPT101 series had a smaller receptive field (~2.5mm), but was primarily sensitive to the infrared range. Since most monitors give off a fair amount of heat, this chip was greatly ineffective. The Agilent QT333 CMOS chip, aside from being beyond frustrating to work with because of its heat sensitivity and QFN case packing, had too large of a receptive field (~5mm) to meet the spatial resolution desired for the device. The Hamamatsu S9706 Digital CMOS sensor gave 12-bit output for red, green, and blue in a serial output array and had the advantage of two resolutions, low (0.36mm x 0.36mm) and high (1.2mm x 1.2mm). Unfortunately, the sensor requires an integration time that is proportional to the log of the luminosity of the light and the log of the number of bits outputted on each channel. Typically, the luminosity for most LCD monitors is between 200 and 500 lumens; using the low value of 200 lumens, the low resolution would require an estimated 25ms integration time to output only 4 of the 12 bits for each color. The S9032-02 chip Analog Color sensor utilizes red, green, and blue color diodes with a total area of 2π mm^2, which puts it in the range of the spatial sensitivity desired for the device. Initially, the device was rejected as having too poor of a spatial resolution, preliminarily estimated to be around 3mm. However, the sensor came with a lens that was not needed for our application; by removing this lens, the sensor obtained a more acceptable spatial

resolution. Ultimately, the S9032-02 RGB color sensor from Hamamatsu was chosen as the sensor for the device, as it met the spatial resolution requirement, the analog signal output adds only negligible latency, it is easier to work with by not requiring several additional timing circuits, and costs less per chip ($5 for the S9032-32 compared to $19 for the S9706).

3.4.1.2: The Haptic Actuator

Due to the limitations and the possible safety issues concerning the use of the pager motor as a haptic actuator, it was abandoned and focus was shifted on finding a new actuator capable of providing more than a single monotonic vibration and with better control of the amplitude. The most intuitive way to produce multi-tonal outputs to synthesize temporally generated textures is to utilize a speaker element, either a voice-coil or piezoelectric crystal, as the haptic actuator. Voice-coils are electromagnetically driven linear actuators consisting of a cylindrical coil of wires surrounding a magnetic core. When current passes through the coil of wires, a magnetic field is generated that either attracts or repels the magnetic core, depending on the direction of the flow of current. Voice-coil speakers have the advantage of being able to be driven by lower voltages than piezoelectric materials, and can produce greater displacements than piezoelectric speakers, with the force produced being proportional to the number of coils, the diameter of those coils, and the current passing through the coils. The piezoelectric speaker is a disk-shaped piezoelectric device that deforms in response to an electric potential across it. Its displacement can be controlled by increasing the voltage across the material, by increasing the size of the disk, or by sandwiching multiple disks together to

create different layers. It has the advantage of being driven by lower currents and having a slimmer profile than voice-coil actuators. Piezoelectric speakers and voice-coils both have about equal cost, which ranges from around a few dollars for small, commercially available units to hundreds of dollars for custom made parts. Both voice coils and piezoelectric devices have no DC response and very limited response for frequencies below a certain range (typically <10Hz).

After testing about twenty different voice-coils and piezoelectric speakers, a bilayer piezoelectric speaker (Taiyo Yuden) was chosen over the others, due to its thin, flat profile (~2mm), which allowed the overall device to be smaller. It also had a stronger vibratory output compared to the commercially available voice-coils of similar diameter that we tested. In comparison to monolayer piezoelectric speakers of equal size, the two layers of the bilayer actuator yielded a greater displacement for a given voltage, while only having a cost of approximately $10. The piezoelectric speaker used has a resonance frequency of 2500 Hz, but can output a signal from 1 Hz to over 20 kHz . The piezoelectric has a perceivable range from 5 Hz to 500 Hz, but the ability to detect small changes in frequency within the lower band-range is fairly poor. However, for faster scanning velocities (100mm/s), the practical frequency limit is 20 Hz or higher. Further, it can sustain a large range of voltages, from $1V_{p-p}$ to over $40V_{p-p}$ without damage. Its thin profile (<2mm thick) allows it to easily fit into compact cases, simplifying the design of the overall device. The main problems with the actuator are that: (1) sufficient pressure applied directly on the surface of the piezoelectric speaker will dampen the vibrations, although this is avoided with the device case design, and (2) it is unknown how well the

ceramic material will hold up against the skin's natural oils, though evidence of degradation has yet to be observed. The piezoelectric speaker is driven directly by the output of the computer's D/A (max 10V, 5mA); although additional amplification of the voltage of the output signal and providing a bit more current to the speaker than the 5mA the DAQ card can output could possibly greatly improve upon the perception of the output. Figure 9 shows an early prototype for a wearable case that would

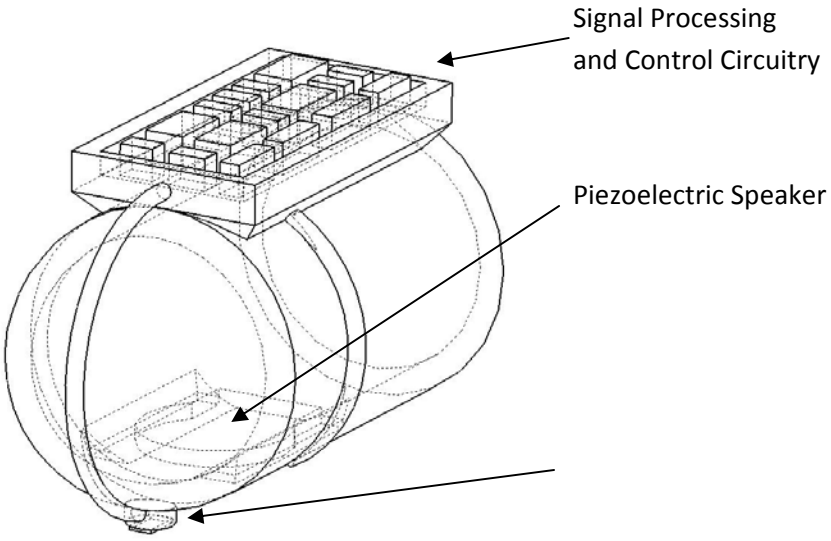

Figure 9 the preliminary design for the third prototype

incorporate the piezoelectric speaker. The idea behind this design was not far from the current design, except the current design calls for: a) some of the circuitry to be placed on the same board as the optical sensor with the rest in a case on the back of the hand or wrist, and b) that the finger-mounted case is worn with an elastic band rather than having a solid case that surrounds the whole fingertip. However, by using the piezoelectric speaker the principle design feature was to place the haptic actuator directly against the skin, between the finger and the optical sensor, for the best possible kinesthetic concordance.

3.4.2: Overall Device Design

As mentioned before, a "glove-like" design was chosen to allow for multiple interaction points. However, the previous prototype, which used a construction glove to serve as a means to mount components, was not very comfortable to wear. The most current prototype eliminates the use of an actual glove while continuing to use something that could be worn. An elastic band is used to mount the case containing the device components to the finger. This has the added advantage of fitting more people than a single glove without the need to make any adjustments. Also, as with the previous device, the prototype also continues to allow for independent device movement for multiple devices on multiple fingers without the user having to hold them. This allows the fingers to have more of a natural, independent exploration of the visual image.

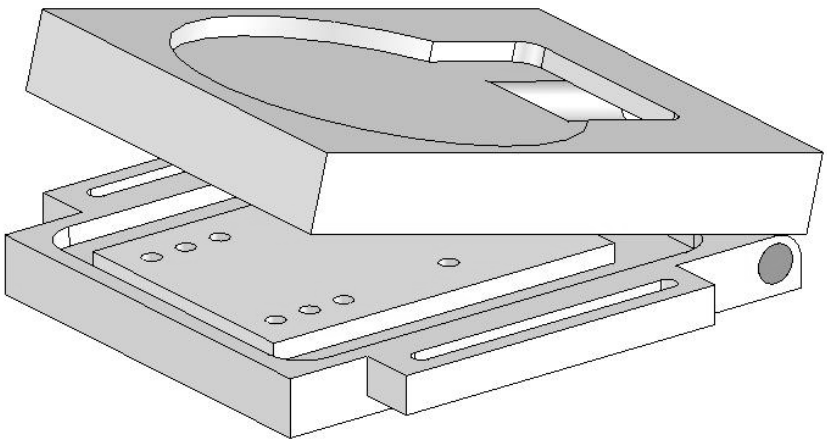

Figure 10 the modified design for the third prototype casing

Figure 10 shows the concept for the third version of the device. The device consists of three layers held together by two screws. The first layer, the bottom part of the case, is the portion of the case that contacts the graphic surface. It has a pinhole aperture through

which the optical sensor gets light through, shown in Figure 11 as the small ~2mm hole in the center. This small pin hold acts to increase the accuracy of the spatial resolution for the optical sensor by restricting light transmitted towards the sensor at non-orthogonal angles. Both are located concentrically underneath where the center of the piezoelectric speaker sits, which is also roughly where the user's fingertip rests on the device. It also contains a small rubber pad that contacts a push-button switch located on the next layer, the PCB (not shown in the picture).

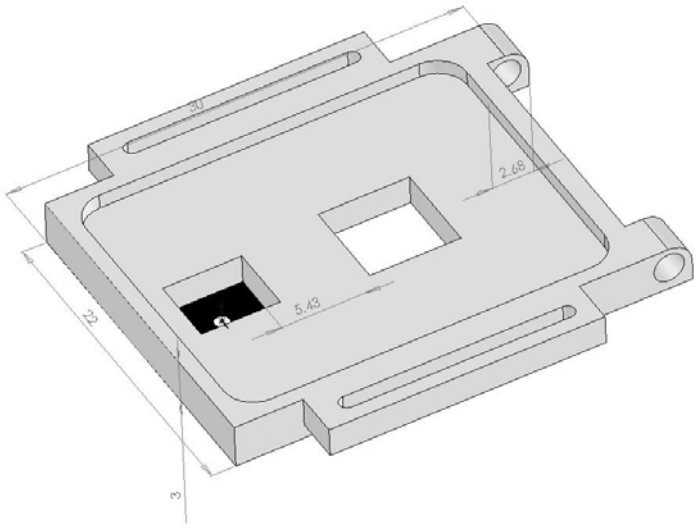

Figure 11 the bottom piece of the third prototype case

Next, the PCB board rests inside the top and bottom case parts and is held in place by the two screws that pass through it. The PCB board contains the RGB color sensor, the pre-DAQ circuitry, and a push-button switch. The push-button switch on the PCB board serves as the on-off button for the device. The device turns on only when the device is pressed against a surface; the weight of the device alone is not a sufficient enough force

to do this, so it requires active pressure by the user. One benefit of this is that the device will only require power when it is being used. A second advantage is that the device does not falsely activate (because it will turn off) when it is pointed towards another light source, such as a lamp.

The circuitry present on the chip provides line buffering, low-pass filtering, and adjustable gain for each of the three color channels. Figure 12 shows the circuit diagram [See Appendix 8.1 for full size copy]. The first set of operational amplifiers serve as a buffer for the photodiode and as a low-pass filter with a cut off frequency of 106Hz. The filtering component was necessary to remove high-frequency noise generated by the photodiode and to band-limit the signal such that the second-stage op-amps could properly increase the signal potential. These inverting amplifiers (the buffer inverted the original signal, so by having an inverted gain the signal was corrected) provided an adjustable gain for each of the channels, so that their outputs could be calibrated. The push-button switch provided an interrupt in the piezoelectric speaker line; however, when the device is driven off an internal power supply such as a battery, the push-button switch will serve as an interrupt between the power supply and the rest of the circuitry.

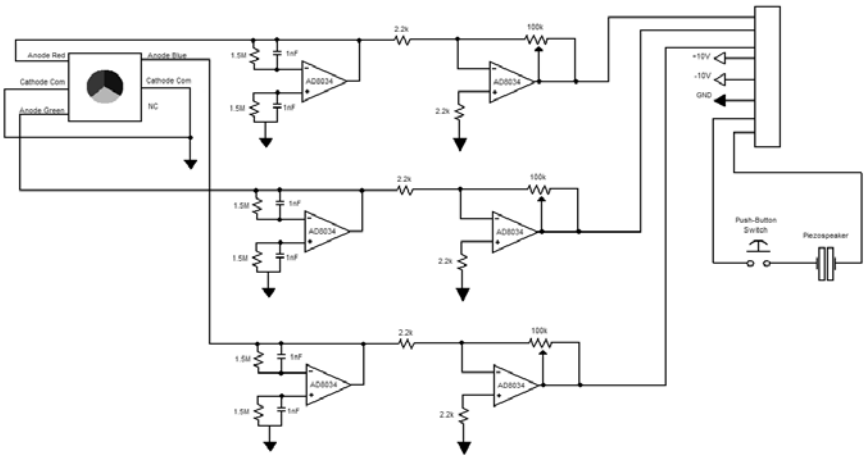

Figure 12 the pre-DAQ circuitry

The top part of the device case houses the piezoelectric speaker, the screw terminals, and serves as the contact surface for the user's finger. An elastic band hooks into a groove located between the bottom and top portions. It holds the finger in place and stretches for many different finger sizes. Additionally, the contact surface for the finger on the device is shaped such that the distal eminence of the fingertip slightly rests on the piezoelectric speaker, but so the remainder of the distal phalanx rests on an inner ledge of the casing. This: 1) provides a better seating of the finger on the device to prevent slippage of the finger in the device as it drags over the surface being explored and 2) prevents the pressure applied by the finger from dampening the vibrations of the piezoelectric speaker. Figure 13 shows the top portion of the prototype case.

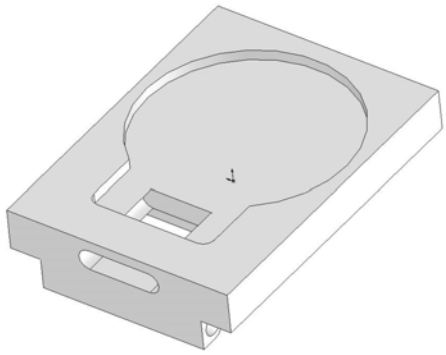

Figure 13 the top part of the third prototype case

3.4.2.1: Control Program Design

The control for the program design stemmed from earlier hardware designed to control the piezoelectric speaker. This hardware consisted of three voltage comparators that used the output of a photodiode. This photodiode was the (later) rejected BCS series device that consisted of only one output and had sensitivity to light in the blue wavelength spectrum. The three comparators acted as a parallel Analog to Digital converter, generating a 2 bit output. This two bit output was then tied to a bank of clock circuits operating at different frequencies; if the binary output was 11, then the logic circuit output one frequency, if 10 then another, if 01 then another, and if 00 there would be no output. Ultimately, this became the archetype for the later software control for the S9032-02 RGB color sensor and piezoelectric speaker.

For the final prototype, using Labview, the three color channels, as shown in the circuit diagram on Figure 12, are sampled with a NI ADC 6230 16-bit PCI DAQ board at a rate of 4 kHz. The program was set up so that while it was running, it would scan through a Case structure containing multiple comparisons of the three input color

channels. Each case comparison was based on the mean voltages for the red, green, and blue channels for a particular color. When a case was chosen it would output a unique texture. To do this the while loop stopped running and the program continuously output the corresponding case tone until the case became false. At that point the while loop would resume scanning through the case structures until one was true, as long as the program was still running. If no case was true, it would default to the last case, which corresponded to the color black that had no output tone. Figure 14 shows the Labview block diagram of the control program.

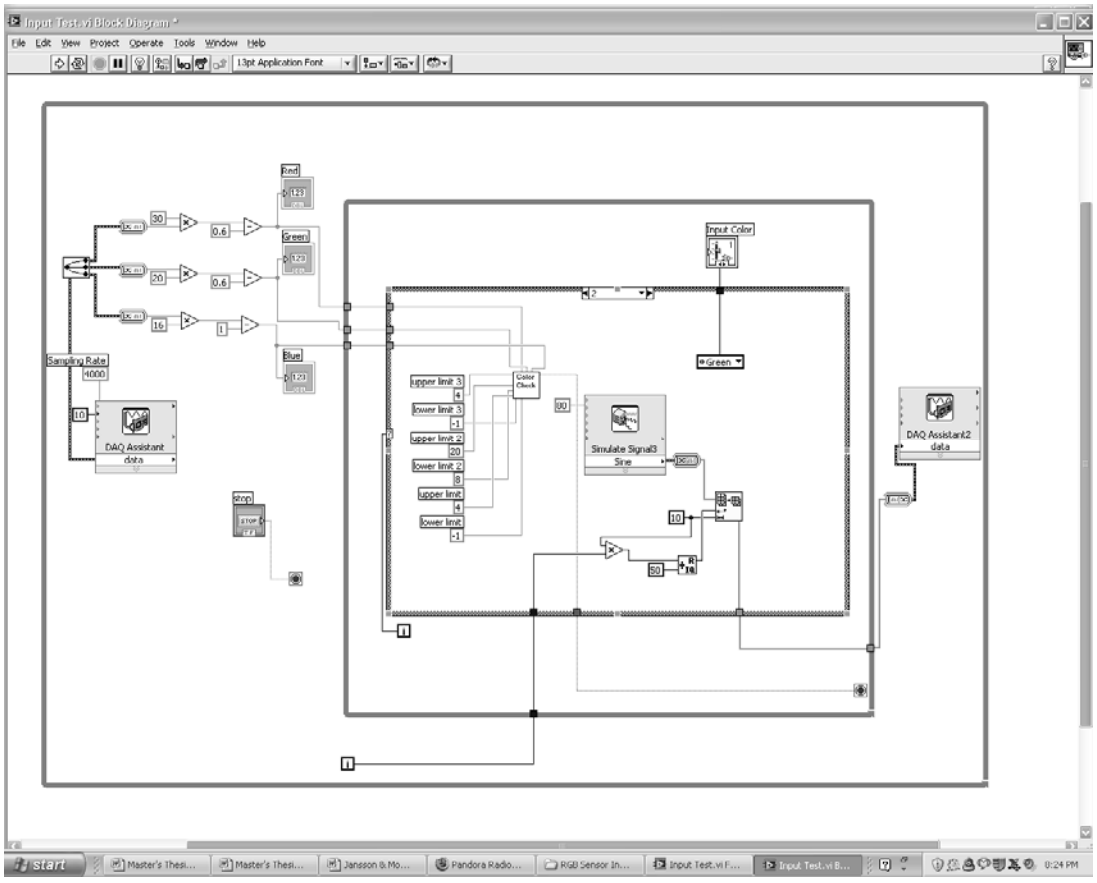

Figure 14 the LabView block diagram

Each tone was generated using the Simulate Signal VI in Labview; in each case it was set to generate 10 samples at 4 kHz and set to run as fast as possible. The output was sent to an NI DAC 6230 16-bit PCI DAQ (the card has channels for both analog input and output) set to generate continuous output synchronized with the Simulate Signal VI. Currently there is a problem with the initialization of the while loop that causes a delay in switching between cases. This software problem will be worked out later or will be made irrelevant by the implementation of an embedded controller.

3.4.2.2: Texture Set

The use of Just Noticeable Differences (JNDs) for amplitude and frequency [14] was found to be inappropriate for generating textures, as it was extremely difficult to perceive these differences. Even van Erp's guidelines [15] of using 4 amplitude levels and 9 frequencies seemed unrealistic in our tests, considering the desire to have subjects use the minimal amount of effort to identify the textures. In addition, amplitude and frequency interact in the perception of roughness [16, 51] eliminating the effective use of one or the other to generate texture sets. Since vibration frequency has a larger range of perceivable options, it was decided to keep output amplitude fixed and use frequency variation for generating texture sets. The texture set chosen consists of a set of three single-note tones that are modified by four different conditions, plus one additional background tone. A tone is generated by the Labview control program through the Simulated Signal Express VI; each tone consists of either a single note (a sinusoidal waveform), a multi-note tone (addition of harmonic notes), or a modulated tone (two notes multiplied together). The set specifically consists of 45Hz, 75Hz, and 145Hz sine waves (the pure notes), the same

frequencies presented as either sawtooth waves or square waves (multi-note tones), or modulated with a 15Hz carrier wave. An additional output of a 400 Hz square-wave was added to be used for borders, if needed. The frequencies were, in part, chosen based on the use of modulated signals from Weisenberger's work [42], which found that higher frequencies modulated with lower frequencies (within the range of 15-30Hz) had a higher performance in discrimination tasks. The second consideration was to avoid using frequencies for which the channels of vibration perception were most sensitive to (i.e. 200-300 Hz for the PC and 25-40 Hz for the NPI channel), as they had the fastest adaptation times [3]. To further reduce the effects of adaptation by reducing the load on the PC channel (i.e. by not having all the textures stimulate the PC channel), the set of textures chosen stimulates both the PC (80-500 Hz) and non-PC channels (3-100 Hz). The final consideration was that waveform and frequency had been successfully used to create larger sets of textures [52, 53]. The waveforms chosen represent the most mathematically dissimilar choices among the superset of sine, cosine, square, triangle, and sawtooth waves. The evaluation of dissimilarity was made based on an evaluation of the additive synthesis of harmonics for each waveform. Sine and cosine waves differ only in phase shift of 90 degrees and so have the highest degree of similarity, so only one choice was used. Square waves contain the odd harmonics of the fundamental frequency divided by a ratio of 1/n, where n is the odd harmonic value, making it the least similar to sine waves. Triangle waves also contain the odd harmonics of the fundamental frequency, but are divided by a ratio of $1/n^2$, which while somewhat dissimilar from square waves, makes the additive contribution of the harmonics less significant than for the square

wave, thus making the overall waveform somewhat similar to the sine wave. The sawtooth wave contains both the odd and even harmonics divided by a ratio of 1/n, making the waveform somewhat dissimilar to both sine and square waves, and the most appropriate choice for the third waveform modifier.

This set was chosen to provide greater perceptual diversity of the textures output by the device, based on previous research by other groups. Each output had a peak-to-peak amplitude of 20V (the maximum allowed output of the NI DAQ card). One method used to "boost" the perception of the output was to compress the metal frame of the piezoelectric speaker inward along its border. This was accomplished by clamping the piezoelectric speaker on either side. To test the effects of clamping the piezoelectric speaker, the acceleration amplitude of the vibration was measured using a FS20 Series Load Cell from MSI Sensors. The load cell was placed underneath the piezoelectric speaker so that the displacement of the speaker would compress the load cell. The results of this test are listed in Appendix 8.3. Clamping the speaker did provide additional deflection in the piezoelectric, with accelerations 2 to 3 times greater with the clamping, possibly by forcing the speaker into a convex shape. However, it also greatly reduced the fidelity of the signal output and greatly distorted the output characteristics of the piezoelectric speaker. This was measured using a second piezoelectric device, a cantilever beam shaped piezoceramic, and an oscilloscope. It was decided that boosting the voltage and current available to the device was a better solution for boosting the output performance of the haptic actuator.

3.4.2.3: Color Set

The device was chosen to be made sensitive to fourteen colors, from which a smaller subset will later be selected and paired with a vibratory texture tone. This was based on the work of Thompson and her colleagues who found that using spatial arrangement and geometric orientation cues within a 2-D representation greatly increases object identification [2, 7]. Their method, TexyForm, utilizes three variations of a single texture to represent vertical, horizontal, and radial directional cues. To incorporate their methods, we decided to make the device selectively sensitive to fourteen colors: three base colors (red, green, and blue) to represent three different object parsing options without any spatial direction cues, nine colors (yellow, purple, aquamarine, and all the half-saturation colors) to represent the options of each base color with a different directional cue, and two colors (black and white) to represent borders and backgrounds for the image. These colors were chosen based on information of using effective color contrast for people with low vision [12].

3.4.3: Device Testing

3.4.3.1: Spatial Collocation

The evaluation of the spatial resolution for the device was determined for two factors, the thinnest line (absolute resolution) and the smallest separation between two lines (relative discriminatory resolution) that the device can detect. To test the spatial resolution, a series of 392 parallel lines with line thicknesses of approximately 0.3mm, 0.6mm, 0.9mm, 1.2mm, 1.5mm, 1.8mm, and 2.1mm were drawn, with each set of line thicknesses having 56 lines total. Each set of 56 lines was then further broken down into

individual sets of 8 lines, with each smaller set of lines having the line separations of approximately 1mm, 1.5mm, 2.0mm, 2.5mm, 3.0mm, 3.5mm, and 4.0mm, with the separation distance measured from the end of one line to the beginning of the next. [See Appendix 8.4.1 for example test].

Next, in order to determine whether the spatial resolution varied with color, the 392 lines were repeated for different combinations of line and background color. The color/background combinations analyzed were: white, red, blue, green, yellow, purple, and aquamarine lines, both at full and half saturation, against a black background; red and green lines against a yellow background, red and blue lines against a purple background, and green and blue lines against an aquamarine background; and finally red, green, blue, yellow, purple, and aquamarine lines of full saturation with their half saturation values as a background.

The color combinations used in the spatial resolution test shows baseline ability of the device to detect differences between the poorest color contrasts available, but not as a guideline of what colors should be used adjacent to each other. An alternative method to encode object orientation is to more directly mimic TexyForm by using a separate color, such as a border color, to draw lines indicating 3-D object orientation over a background color. This method requires no modification in the device design, and thus was not a factor in the design, even though the method may be tested in the future.

All combinations were drawn using MSPaint; saturation values were manipulated by adjusting the RGB value for the color within the program. Full color saturation refers

to a value of 255 for a particular color (e.g. yellow would have a RGB value of 255, 255, 0), while half saturation refers to a value of 128, or half the full saturation level. The ability to detect the correct color was assessed on a true or false basis: the device had to correctly identify the color for every line in a set to be considered true. If the device was able to correctly identify most of the lines (five or more of the eight in the set), but not all, then the set was repeated on the basis that human error in scanning the lines might have occurred. Upon repeating the trial, if the device still could not determine the color than a false statement was given for that test. [See Appendix 8.4.2 for results].

It was found that against a black background, the full-saturation colors (red, green, blue, yellow, purple, and aquamarine) have an absolute spatial resolution of 1.2mm and a relative discriminatory resolution of 1.5mm. The corresponding half-saturation colors have a poorer absolute resolution at 1.8mm, but a better relative discriminatory resolution at 1mm. For color combinations with poor saturation contrast, i.e. full-saturation colors against a background of their half-saturation colors, the absolute resolution was 0.9mm and the relative discriminatory resolution was 2mm. For color combinations with poor hue contrast, i.e. colors of equal saturation values that would be adjacent on a color diagram, the absolute resolution is 1.8mm and the relative discriminatory resolution is 1.5mm for full-saturation combinations and 2mm for the half-saturation combinations. Thus, in the worst case scenarios, the device still produces reasonable spatial resolution of 2mm.

3.4.3.2 Temporal Collocation

The temporal collocation of the device remains an issue, due to a problem with the software control program. Initially, the program was set to output 20 samples at a rate of 10 kHz. However, after the samples were output, the DAQ would reinitialize, creating a 1ms delay. This would not only increase the latency of the output by 1ms, to a total of around 10ms (an acceptable amount), it would dramatically alter the output signal waveforms (not acceptable). The DAQ output was changed to output continuously and synchronized to the Simulate Signal VI; this eliminated the DAQ re-initialization. Unfortunately, shortly after this correction, it was noticed that the program had a variable delay when selecting a case that ranged from 400ms to over 800ms—far from acceptable. The most likely source of the error is in the calculation of the output waveforms; however, this is unknown as of this time, as the problem could have to do with computer or DAQ card. Ultimately, the goal is to have the device running with less than 10ms of signal latency between light detection by the photodiode and the piezoelectric speaker activating.

3.4.3.3 Actuator Feedback

Since the ability to generate synthetic textures is a device feature aimed to greatly improve the perception of the 2-D graphics, it is necessary to look at the prominence of the textures chosen. As previously mentioned, the Labview program was set up to identify and output 14 possible textures, but that a more realistic number of usable textures will be around 8 to 9. The biggest concern was not only to see which outputs were most salient and which were most confused, but to see if there were any trends to

this confusion within the texture sets. Six sighted subjects with no visual impairment between the ages of 22 and 32 years were tested with the device to see how well they could distinguish the 14 different output modes of the device (no texture is considered a possible texture output choice for the test). Digitally filtered pink noise was presented during both phases to eliminate auditory feedback. Each subject went through a training period consisting of a brief exposure and description of each output choice, followed by a practice exercise. During the practice exercise, subjects had to try and guess the frequency (low, medium, or high) and waveform type of the output signal (sine wave, square wave, sawtooth wave, and modulated wave), or whether the texture was the 400 Hz border texture or no texture, and they received feedback on their choices. Subjects could also ask to go back and feel specific textures of their choosing. The test resumed after a 5 minute rest; subjects were presented the 14 output types in random order and asked to identify them as in the practice exercise. As with the practice exercise, subjects could request to go back and sample any two specific texture outputs, but were given no feedback or other help in their choices to select.

G/C	0	1	2	3	4	5	6	7	8	9	10	11	12	13
0	6										2	1		
1		3							1					
2		1	3			1		1						
3			1	2		1			2	1				
4			1		2			1				1		
5			1			2								
6				2			5	1		2				
7					1			2						
8		1				2			3					
9				2			1			3				

10		1			1			1			4	1	1	
11					2							4	1	
12													4	
13														6

Table 2 the confusion matrix for the texture sensing test

Table 2 shows the confusion matrix for the test; the columns are the correct textures and the rows are the guessed textures. Subjects were not told that each texture was presented only once, and thus each row does not necessarily add up to six, whereas each column does. The color coding on the table helps show the confusion trends between low (light gray), medium (medium gray), and high (dark gray) frequencies of each waveform type; numbers 1-3 correspond to the sine waves, 4-6 correspond to the sawtooth waves, 7-9 correspond to the square waves, 10-12 correspond to the modulated waves, 0 is the 400 Hz border frequency, and 13 is no texture. As it can be seen by the table, the highest accuracy (100%) was in detecting the 400 Hz border frequency and no texture. The low and medium modulated textures were confused with the 400 Hz border 33% and 17% of the time. Modulated signals were also confused with different modulated signals some of the time, with no apparent specificity. The greatest confusion was between sine-waves, sawtooth waves, and square waves of the same frequency, the exception being the high frequency sawtooth wave. A possible solution would be to eliminate either the sawtooth waves or the square waves; however, a test with more subjects is needed to perform a comparison test with enough power to be of use. [See Appendix 8.2 for data]

3.4.3.4: User Safety

International Standards ISO 5239-1 and ISO 5349-1 provide a guide for the maximum daily vibration exposure for an individual by proposing a frequency-weighted algorithm for measuring the acceleration magnitude of these vibrations. However, Dong and his colleagues found that the energy absorption distribution for vibration is dependent on the frequency, with higher frequency (>100Hz) energy being absorbed more locally. Further, they suggest that for frequencies higher than 16Hz, with some exceptions, "the relative weighting is higher than that of the ISO" [9]. All of the outputs for our device fall within this frequency range; however, only one frequency was chosen that was much higher than 100 Hz, the 400 Hz frequency for borders, and it was never designed to be a texture for use over large areas of the image. Furthermore, the power output for the piezoelectric speaker is so small compared to the mass of the finger that the average power absorption density should be low enough with these guidelines where even temporary symptoms of vibration disorders are not a concern.

3.4.3.5: Adaptation

As discussed with the previous prototype design, adaptation is expected to be an issue over time, making the device frustrating to use. To somewhat limit the affects of adaptation, sets of temporally-encode textures were chosen that stimulate both PC and non-PC channels, so that every texture chosen does not stimulate just the PC channel (which is sometimes seen in devices) due it its high sensitivity. However, some adaptation is still expected, as haptic exploration of multiple images can be a time consuming task.

3.4.4 Discussion

Overall, the third prototype accomplished the main goal of using color detection to encode a variety of textures that encoded both object parsing and 3-D orientation cues about an image. However, in testing the actuator output of the textures, it was determined that greater attention to the texture design parameters was needed, and that prior results from other groups (namely [49, 50, 52]) are not applicable to the current design setup being tested. The characteristics of the optical sensor work very well: it has satisfactory spatial resolution (less than 2mm for all tests) and can accurately sense a large number of colors. The characteristics of the piezoelectric speaker are also adequate: it can have a fairly strong output while remaining below the acceleration amplitude that causes HAVS-like symptoms. The greatest issue remaining with the device design is the choice of textures and the implementation of the device in real-time. However, these problems do not necessarily require a redesign of the prototype, but rather further testing of the choice of textures and the software control program.

4. Conclusion and Future Work

The final prototype developed for this project succeeded in meeting most of the criteria set out for it; specifically, it is designed for accurate haptic perception of 2-D graphic information through the use of color-encoded, temporally-generated texture feedback distributed over multiple fingers as independent, point-contact display devices. The range of textures capable of being produced by the device has the potential to allow for proper user discrimination and can be used to encode both object segmentation and 3-D orientation. This will enable the device to mimic the function of TexyForm, a very successful means of creating perceivable static tactile images. By having the capability for expansion to allow for the use of independent devices on multiple fingers, the design potentially enables parallel haptic processing of the textures, facilitating haptic exploration by reducing the time needed to explore the image. Through using graphic colors as a means to encode the different texture outputs, images are easier to create into a usable form while retaining important visual information that can still be accessed by individuals who have low vision. All of this function comes in a device costing less than $100 that is both safe and easy to use.

Future work will focus around developing an efficient and highly perceivable set of textures to encode object parsing and 3-D orientation, as well as the development of multiple finger models of the device. Particularly, more time will be spent on testing which parameters are functionally perceivable for the intended use of the device (compared to what parameters can be slightly discriminated in threshold tests) by

performing multiple comparisons tests and seeing which textures have the greatest contrast. Once a set of textures has been chosen, then work on the multiple finger design will begin.

5. Support

This work was supported in part by NSF IIS grant #0712936 and an A.D. Williams grant.

6. References

1. Krufka, S.E. and Barner K.E. (2005). "Automatic Production of Tactile Graphics from Scalable Vector Graphics". *Assets '05*, October 9-12, Baltimore, Maryland.
2. Thompson, L.J., Chronicle, E.P., Collins, A.F. (2003). "The Role of Pictorial Convention in Haptic Picture Perception". Perception, **32** (7), pp. 887-893.
3. Jones, L. and Lederman, S.J. (2006). Human Hand Function. Oxford University Press, New York.
4. Johnson K. O., Yoshioka T., and Vega-Bermudez, F. (2000) "Tactile Functions of mechanoreceptive afferents innervating the hand. In Human Hand Function. Eds L. Jones and S. Lederman. Oxford University Press, New York.
5. Weinstein S. (1968) "Intensive and extensive aspects of tactile sensitivity as a function of body part, sex, and laterality." In Human Hand Function. Eds L. Jones and S. Lederman. Oxford University Press, New York.
6. Vega-Bermudez, F. and Johnson, K. O. (2001) "Differences in Spatial Acuity between Digits". In Human Hand Function. Eds L. Jones and S. Lederman. Oxford University Press, New York.
7. Thompson, L.J., Chronicle, E.P., Collins, A.F. (2006). "Enhancing 2-D Tactile Picture Design from Knowledge of 3-D Object Recognition". European Psychologist, **11**(2), Hogrefe and Huber Publishers, pp 1-9.
8. International Standard ISO 5349-1. (2001). "Mechanical vibration – Measurement and evaluation of human exposure to hand-transmitted vibration. Part 1: General Requirements". Switzerland.
9. R.G. Dong, J.Z. Wu, and D.E. Welcome. (2005) "Recent Advances in Biodynamics of Human Hand-Arm System." Industrial Health, **43**, 449-471.
10. Lederman, S.J., Loomis, J.M., and Williams, D.A. (1982). "The Role of Vibration in the Tactual Perception of Roughness". Perception and Psychophysics. **32** (20), pp.109-116.
11. Edman, P.K. (1992) Tactile Graphics. American Foundation for the Blind. New York.
12. A. Arditi. (2008) "Effective Color Contrast: Designing for People with Partial Sight and Color Deficiencies." Lighthouse International. http://www.lighthouse.org/accessibility/effective-color-contrast.
13. U.S. Census Bureau. Americans with Disabilities. 2002.
14. Sherrick and Cholewiak. (1986) "Cutaneous Sensitivity". In Eds. K. R. B.L. Kufman and J.P. Thomas, Handbook of Perception and Human Performance. New York: Wiley.
15. J.B.F. van Erp. (2002) "Guidelines for the Use of Vibro-Tactile Displays in Human Computer Interaction". Eurohaptics.

16. K.U. Kyung and D.S. Kwon. (2006) "Multi-sensory Perception of Roughness: Empirical Study on Effects of Vibrotactile Feedback and Auditory Feedback in Texture Perception". ICAT, 406-415.
17. Levesque, V. (2005). "Blindness, Technology and Haptics". Technical Report TR-CIM- 05.08. McGill University, Canada.
18. Sjostrom, C. (2002). "Non-visual Haptic Interaction Design". Ph.D Thesis. Lund Institute of Technology.
19. Kurze, M. (1996). "TDraw: A Computer- based Tactile Drawing Tool for Blind People". *Assests '96*. Vancouver, Canada.
20. Kurze, M. (1999). "TGuide: a Guidance System for Tactile Image Exploration". Behavior and Information Technology, **18**(1), pp. 11-17.
21. Hollins, M., Bensmaïa, S., Karlof, K., and Young, F. (2000) "Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling." Perception and Psychophysics, **62**, 1534-1544.
22. LaMotte, R. H. and Srinivasan, M. A. (1991) "Surface microgeometry: Tactile perception and neural encoding." In Human Hand Function. Eds L. Jones and S. Lederman. Oxford University Press, New York.
23. Lamb, G. D. (1983) "Tactile discrimination of textured surfaces: Psychophysical performance measurements in humans." In Human Hand Function. Eds L. Jones and S. Lederman. Oxford University Press, New York.
24. Berla, E. P. and Murr, M. J. (1975) "Psychophysical functions for active tactual discrimination of line width by blind children." Perception and Psychophysics, **17** (6), 607-612.
25. Tan, H. Z., Reed, C. M., Delhorne, L. A., Durlach, N. I., and Wan, N. (2003) "Temporal masking of multidimensional tactual stimuli." Journal of the Acoustical Society of America, **114** (6), 3295-3308.
26. Israr, A., Tan, H. Z., and Reed, C. M. (2006) "Frequency and amplitude discrimination along kinesthetic-cutaneous continuum in the presence of masking stimuli." Journal of the Acoustical Society of America, **120** (5), 2789-2800.
27. Fritz, J.P. and Barner, K.E. (1999). "Design of a Haptic Data Visualization System for People with Visual Impairments". IEEE Transactions on Rehabilitation Engineering, **7** (3), 372-384.
28. Kirkpatrick, A. E., & Douglas, S. A. (2002). "Application-based Evaluation of Haptic Interfaces". *Proceedings of the 10th Symposium On Haptic Interfaces For Virtual Environments and Teleoperator Systems*. IEEE Computer Society.
29. Yu, W. and Brewster, S. A. (2002). " Comparing Two Haptic Interfaces for Multimodal Graph Rendering". *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 3-9.

30. Wies, E.F., Gardner, J.A., O'Modhrain, S., Hasser, C.J., Bulatov, V.L. (2001). "Web-Based Touch Display for Accessible Science Education". Lecture Notes in Computing Science, 2058, pp. 52-60.
31. Roth, P., Kamel, H., Petrucci, L. and Pun, T. (2002). "Do You Feel What I Hear?" Journal of Visual Impairment and Blindness, **96**.

32. Roth, P., Richoz, D. and Pun, T. (2003). "Multimodal System for the Non-visual Exploration of Digital Pictures". *Interact 2003, 9th ICIPTC13 Int. Conf. on Human-Computer Interaction*, Zuerich, Switzerland, Sept. 1-5, 2003.

33. Lederman, S.J. and Klatzky, R.L. (2004). "Haptic Identification of Common Objects: Effects of Constraining the Manual Exploration Process." Perception and Psychophysics, **66**(4), pp. 618-628.

34. Klatzky, R.L., Loomis, J.M., Lederman, S.J., Wake, H., and Fujita, H. (1993). "Haptic Identification of Objects and Their Depictions". Perception and Psychophysics. **54**(2), pp. 170-178.
35. Bouzit, M., Burdea, G., Popescu, G. and Boian, R. ,(2002). "The Rutgers Master II-New Design Force-Feedback Glove". IEEE/ASME Transactions on Mechatronics. **7**(2), pp. 256-263.
36. Lederman, S.J. and Klatzky, R.L. (1997). "Relative Availability of Surface and Object Properties During Early Haptic Processing". Journal of Experimental Psychology: Human Perception and Performance, **23** (6), pp.1680-1707.
37. Lederman, S.J. and Klatzky, R.L. (1997). "Designing Haptic Interfaces For Teleoperational and Virtual Environments: Should Spatially Distributed Forces Be Displayed To The Fingertip?" *Proceedings of the ASME Dynamic Systems and Control Division*, **DSC-Vol 61**, pp. 11-15.
38. Vidal-Verdu, F. and Hafez, M. (2007). "Graphical Tactile Displays for Visually-Impaired People". IEEE Transactions on Neural Systems and Rehabilitation Engineering. **15** (1), pp. 119-130.
39. Pasquero, J. (2006). "Survey on Communication through Touch". Technical Report TR-CIM 06.04. Mc Gill University, Canada.

40. Jeong, W. and Gluck, M. (2002) "Multimodal Bivariate Thematic Maps with Auditory and Haptic Display." Referenced by Plaisant, C., in "Information Visualization and the Challenge of Universal Usability." *Exploring Geovisualization*, Elsevier, 2005.
41. Van den Doel, K. (2003) "SoundView: Sensing Color Images by Kinesthetic Audio." *Proceedings of the 2003 International on Auditory Displays*, Boston, MA.
42. Weisenberger, J. M. (1985) "Sensitivity to amplitude-modulated vibrotactile signals." Journal of the Acoustical Society of America, **80**(6), pp 1707-1715.

43. Loomis, J.M. (1979) “An Investigation in Tactile Hyperacuity.” Sensory Processes, vol **3**, pp 289-302.
44. Verrillo, R.T. and Gescheider, G.A. (1975) “Enhancement and Summation in the Perception of Two Successive Vibrotactile Stimuli. Perception and Psychophysics, **18**, pp 128-136.
45. Hollins, M., Lorenz, F., and Harper, D. (2006) “Somatosensory Coding of Roughness: The Effect of Texture Adaptation in Direct and Indirect Touch.” Journal of Neuroscience, **26**(20), pp 5582-5588.
46. Hollins, M. (2007) “The Coding of Roughness.” Canadian Journal of Experimental Psychology, 61(3):184-95.
47. Craig, J. C. (1985) “Attending to two fingers: Two hands are better than one.” Perception and Psychophysics, **38**(6), pp 496-511.
48. Craig, J. C. and Qian, X. (1997) “Tactile pattern perception by two fingers: temporal interference and response competion.” Perception and Psychophysics, **59**(2), pp 252-265.
49. Brown, L. M., Brewster, S. A., and Purchase, H. C. (2005) “A First Investigation into the Effectiveness of Tactons.” *Proceedings of World Haptics 2005*, Pisa, Italy. IEEE Press, pp 167-176.
50. Brewster, S. A. and Brown, L. M. (2004) “Tactons: Structured Tactile Messages for Non-Visual Information Display.” *Proceedings of AUIC 2004*, Dunedin, New Zealand: Australian Computer Society, pp 15-23.
51. Taylor, B. (1977). “Dimensional interactions in vibrotactile information processing.” Perception & Psychophysics, **21**(5), pp 477-481.
52. Hoggan, E. and Brewster, S. A. (2007) “New Parameters for Tacton Design.” *Conference on Human Factors in Computing Systems 2007*, San Jose, USA. CHI 07 Electronic Publications.
53. Ternes, D. and MacLean, K. E. (2008) “Designing Large Sets of Haptic Icons with Rhythm.” *EuroHaptics 2008*, Springer-Verlag, pp 199-208.
54. Magee, L. E. and Kennedy, J. M. (1980) “Exploring Pictures Tactually.” Nature, **283**, pp 287-288.
55. D’Anguilli, A., Kennedy, J. M., and Heller, M. A. (1998) “Blind children recognizing tactile pictures respond like sighted children given guidance in exploration.” Scandinavian Journal of Psychology, **39**(3), pp 187-190.
56. Kurze, M. (1997) “Rendering drawings for interactive haptic perception.” *Conference on Human Factors in Computing Systems 1997*, Atlanta, USA. CHI 97 Electronic Publications.

7. Appendices

7.1 Prototype 2.0 PCB Circuit Diagram

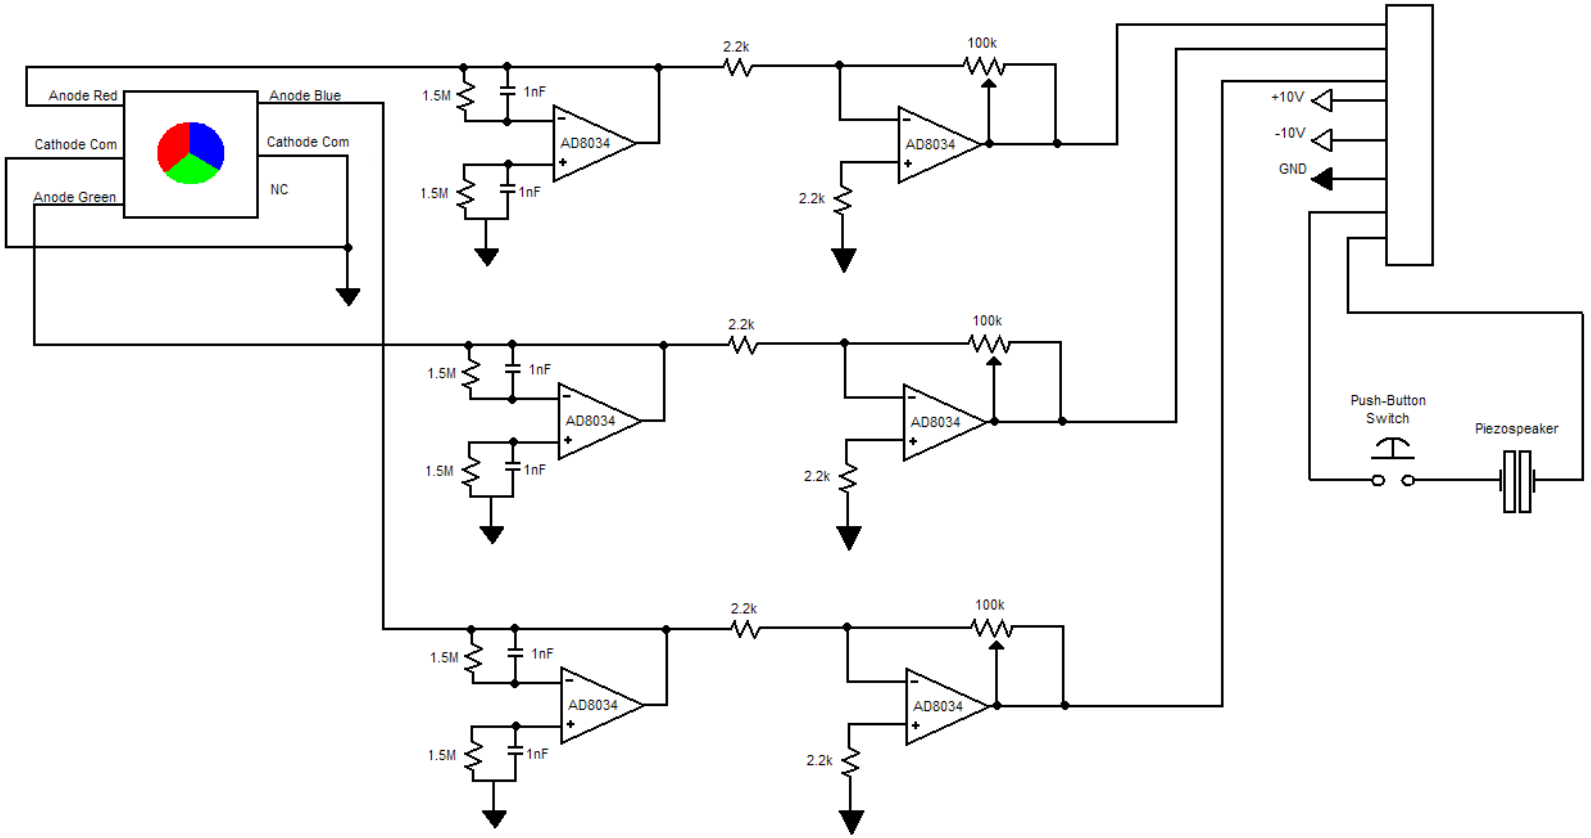

7.2: Haptic Actuator Test Data

Key:

Number	Frequency	Color
0	330 Square	White
1	45 sine	Red
2	80 sine	Green
3	130 sine	Blue
4	45 sawtooth	Yellow
5	80 sawtooth	Purple
6	130 sawtooth	Aquamarine
7	45 square	Dark Red
8	80 square	Dark Green
9	130 square	Dark Blue
10	45*15 sine	Dark Yellow
11	80*15 sine	Dark Purple
12	130*15 sine	Dark Aqua
13	0	Black

Subject 1		
Number	Guessed	Correct
10	dark yellow	dark yellow
1	red	red
0	white	white
7	green	dark red
9	aqua	dark blue
2	green	green
11	yellow	dark purple
8	dark green	dark green
13	black	black
5	dark green	purple
4	yellow	yellow
12	dark aqua	dark aqua
3	dark blue	blue
6	aqua	aqua

Subject 2		
Number	Guessed	Correct
13	black	black
2	blue	green
1	dark green	red
9	blue	dark blue
4	yellow	yellow
8	blue	dark green
5	purple	purple
6	aqua	aqua
0	white	white
11	dark yellow	dark purple
12	dark purple	dark aqua
3	aqua	blue
10	dark yellow	dark yellow
7	dark green	dark red

Subject 3		
Number	Guessed	Correct
7	yellow	dark red
1	dark yellow	red
9	dark blue	dark blue
6	aqua	aqua
12	dark red	dark red
2	green	green
10	dark yellow	dark yellow
13	black	Black
0	white	White
11	red	dark purple
8	green	dark green
5	dark green	Purple
4	dark purple	Yellow
3	blue	Blue

Subject 4		
Number	Guessed	Correct
10	red	dark yellow
13	black	black
5	purple	purple
8	dark green	dark green
12	dark aqua	dark aqua
1	red	red
9	dark blue	dark blue
4	dark red	yellow
11	dark yellow	dark purple
7	dark red	dark red
3	blue	blue
6	aqua	aqua
2	green	green
0	white	white

6	aqua	aqua
11	dark purple	dark purple
0	white	white
3	aqua	blue
5	green	purple
7	dark yellow	dark red
10	dark yellow	dark yellow
2	yellow	green
12	dark purple	dark aqua
9	dark blue	dark blue
4	dark yellow	yellow
8	dark green	dark green
1	red	red
13	black	black

Subject 5		
Number	Guessed	Correct
13	black	black
4	dark purple	yellow
0	white	white
12	dark aqua	dark aqua
3	dark blue	blue
11	dark purple	dark purple
7	dark red	dark red
5	blue	purple
1	green	red
10	red	dark yellow
8	blue	dark green
6	dark blue	aqua
2	purple	green
9	aqua	dark blue

Subject 6		
Number	Guessed	Correct

7.3: Frequency Response of the Piezoelectric Speaker

7.3.1: Unclamped Piezo-speaker

Frequency	Amplitude (mv)	Frequency	Amplitude (mV)	Frequency	Amplitude (mV)
10	3.43	220	16.7	430	36.5
20	4.18	230	26.8	440	31.8
30	4.50	240	22.8	450	25.0
40	4.12	250	12.5	460	24.3
50	4.18	260	7.5	470	25.6
60	4.18	270	6.56	480	25.6
70	4.18	280	7.81	490	22.8
80	3.68	290	10.6	500	22.8
90	3.68	300	15.6	510	19.3
100	4.00	310	22.1	520	18.7
110	4.06	320	18.7	530	18.1
120	4.08	330	19.0	540	18.7
130	4.18	340	25.0	550	20.9
140	4.25	350	27.1	560	20.9
150	4.25	360	23.7	570	19.3
160	4.18	370	25.6	580	17.5
170	4.31	380	31.8	590	21.8
180	4.62	390	34.0	600	26.8
190	5.06	400	35.0	650	41.2
200	7.62	410	35.9	700	35.6
210	12.1	420	37.5	750	77.5

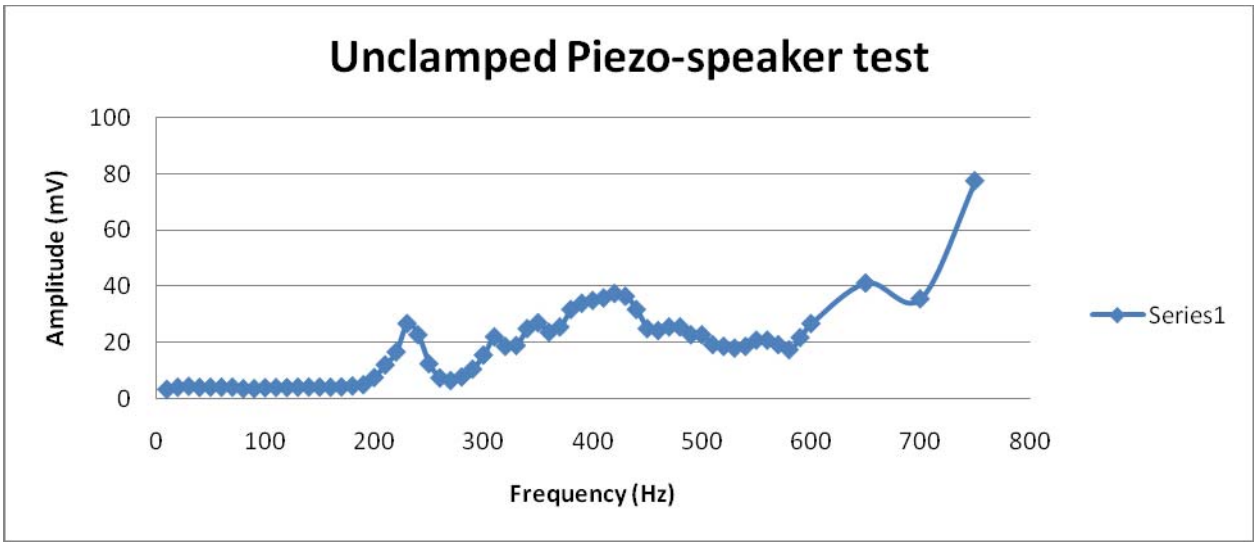

7.3.2: Clamped Piezo-speaker

Frequency	Amplitude (mV)	Frequency	Amplitude (mV)
1	156	270	107
2	150	280	107
3	150	290	107
4	145	300	107
5	132	310	107
6	125	320	107
7	125	330	107
8	123	340	107
9	123	350	107
10	123	360	100
15	115	370	100
20	105	380	100
25	96.2	390	107
30	100	400	115
35	100	410	93.8
40	110	420	93.8
45	127	430	93.8
50	127	440	93.8
60	126	450	112
70	124	460	112
80	124	470	118
90	124	480	118
100	124	490	118
110	125	500	118
120	115	550	106
130	107	600	96.9
140	107	650	96.9
150	107	700	96.9
160	107	750	96.9
170	107	800	96.9
180	107	850	82.8
190	123	900	82.8
200	114	950	76.6
210	109	1000	76.6
220	109	1050	76.6
230	89.3	1100	76.6
240	107	1150	76.6
250	107	1200	76.6
260	107		

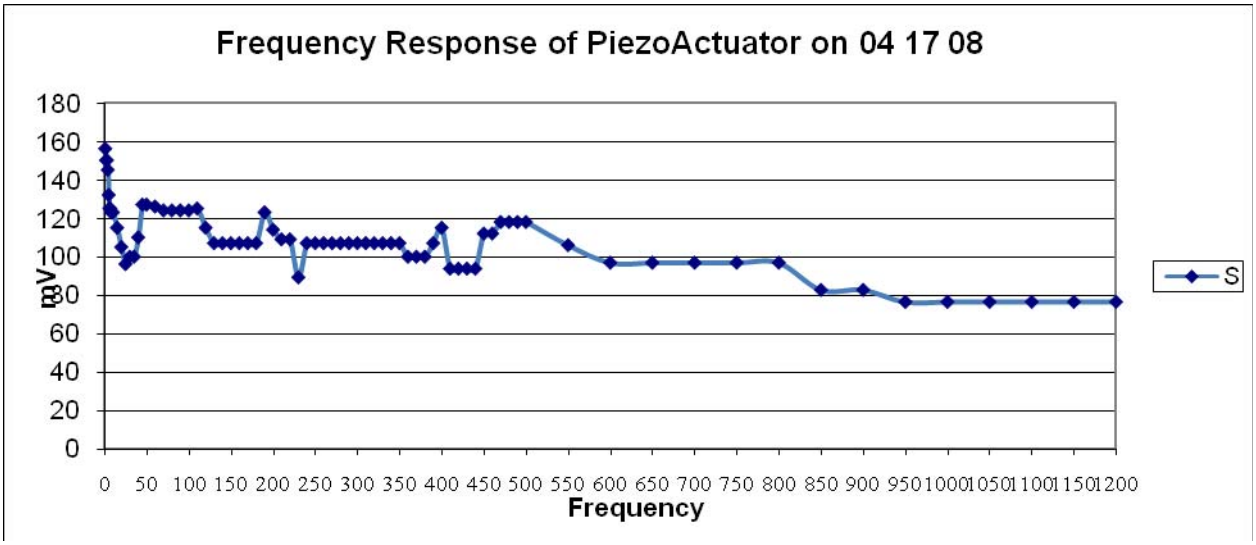
Frequency Response of PiezoActuator on 04 17 08
180
160
140
120
100
80
60
40
20
0
mV
0 50 100 150 200 250 300 350 400 450 500 550 600 650 700 750 800 850 900 950 1000 1050 1100 1150 1200
Frequency
S

7.4 Line Resolution Tests

7.4.1 Test Image

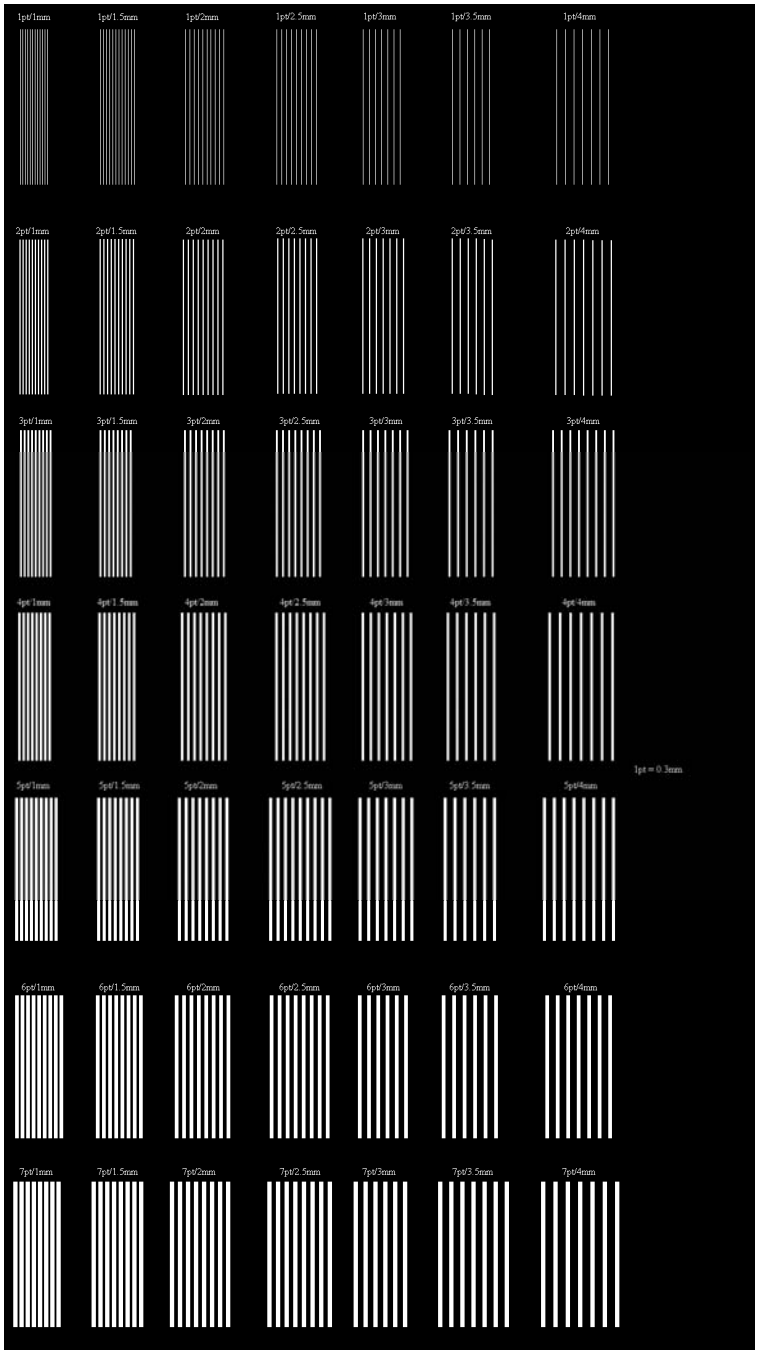

The colors shown (white lines/ black background) were changed to match the test shown.

Image is not to scale.

7.4.2 Test Results

Key: Line Color/Background Color; A = Aquamarine, B = Blue, DA = Dark Aquamarine, DB = Dark Blue, DG = Dark Green, DP = Dark Purple, DR = Dark Red, DY = Dark Yellow, G = Green, P = Purple, R = Red, Y = Yellow.

RGB Color values for the test:

Color	Red Value	Green Value	Blue Value
Red	255	0	0
Green	0	255	0
Blue	0	0	255
Yellow	255	255	0
Purple	255	0	255
Aquamarine	0	255	255
Dark Red	128	0	0
Dark Green	0	128	0
Dark Blue	0	0	128
Dark Yellow	128	128	0
Dark Purple	128	0	128
Dark Aquamarine	0	128	128
White	255	255	255
Black	0	0	0

1pt = 0.3mm (1 pixel)

G/A	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	TRUE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE

2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

R/Y	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	TRUE	TRUE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE

R/P	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

B/P	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	TRUE	TRUE	TRUE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

B/A	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE

G/Y	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

DG/DA	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE

DR/DY	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

DR/DP	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

DB/DP	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

DB/DA	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	TRUE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	FALSE	FALSE	TRUE	TRUE

DG/DY	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

R/DR	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

G/DG	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

B/DB	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

Y/DY	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

P/DP	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE

1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

A/DA	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE

All colors in this part were tested against a solid black background.

WHITE	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	TRUE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

RED	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	TRUE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

GREEN	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

BLUE	1pt	2pt	3pt	4pt	5pt	6pt	7pt

1mm	FALSE	FALSE	TRUE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

Yellow	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

Purple	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE

Aqua	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE	FALSE
1.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
2.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
3.5mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE
4mm	FALSE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE

Dark Red	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

Dark Grn	1pt	2pt	3pt	4pt	5pt	6pt	7pt

1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

Dark Blue	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

D Yellow	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

D Purple	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE

D Aqua	1pt	2pt	3pt	4pt	5pt	6pt	7pt
1mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
1.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
2.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
3.5mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE
4mm	FALSE	FALSE	FALSE	FALSE	FALSE	TRUE	TRUE